



Article

Delayed Fractional-Order Graph Dynamics for Cascade Escalation and Reconfiguration Failure in Integrated Modular Avionics

Oleksandr Korchenko ¹, Olga Torstensson ^{2,*}, Yuliia Kovalenko ³, Dmytro Prokopovych-Tkachenko ^{4,5,6,*}, Oleh Poplavskiy ⁷ and Yevhen Volkov ⁶

¹ Department of Computer Engineering and Cybersecurity, University of the National Education Commission in Krakow, ul. Podchorążych 2, 30-084 Krakow, Poland; oleksandr.korchenko@up.krakow.pl

² School of Information Technology, Halmstad University, 301 18 Halmstad, Sweden

³ Department of Artificial Intelligence, State University of Information and Communication Technologies, 7 Solomyanska Str., 03110 Kyiv, Ukraine; u.kovalenko@duikt.edu.ua

⁴ Department of Cybersecurity and Information Technologies, University of Customs and Finance, 49000 Dnipro, Ukraine

⁵ State Scientific Institution "Institute of Information, Security and Law of the National Academy of Legal Sciences of Ukraine", 01024 Kyiv, Ukraine

⁶ Department of Cybersecurity Systems and Technologies, State University of Information and Communication Technologies, 7 Solomyanska Str., 03110 Kyiv, Ukraine; udgin33v@umsf.edu.ua

⁷ Department of Military Training, University of Customs and Finance, 4900 Dnipro, Ukraine; poplavskiy@umsf.edu.ua

* Correspondence: olga.torstensson@hh.se (O.T.); omega2417@umsf.edu.ua (D.P.-T.); Tel.: +46-738191724 (O.T.); +380-675602397 (D.P.-T.)

Abstract

Integrated Modular Avionics (IMA) integrates safety-critical functions on shared computing and network resources, creating coupling channels through which a local fault may escalate into a catastrophic system-level scenario. This study develops a graph-based fractional-order model for cascade escalation in IMA architectures with communication delays and reconfiguration failures. The architecture is represented as a weighted directed graph of core processing modules, network switches, and remote data concentrators, where each node carries functional degradation and queue-backlog states. The proposed delayed Caputo fractional-order dynamics incorporate degradation propagation, backlog spillover, mixed-criticality priority conflict, and a state-dependent reconfiguration-failure mechanism. We establish well-posedness and positive invariance of the feasible state domain, derive a sufficient cascade threshold that separates a delay-independent, globally Mittag–Leffler stable nominal regime from a supercritical regime in which bistability and catastrophic attractors may occur, and characterize delay-induced oscillatory instability together with a memory-stabilization effect. Numerical experiments on a synthetic 22-node IMA configuration show fault absorption below the threshold, reconfiguration-contained cascades under sufficient supervisory capacity, and global escalation when reconfiguration collapses under load. The results indicate that backlog growth is an early warning signal and that maintaining the cascade threshold below unity while provisioning reconfiguration capacity above the tipping point can support safer reconfiguration-policy design in certifiable avionics.

Keywords: fractional-order dynamics; fractional-order modeling; dynamical systems; integrated modular avionics; Caputo derivative; time delay; cascade failure; mixed criticality; Mittag–Leffler stability; reconfiguration



Academic Editor: Costas Psychalinos

Received: 10 July 2026

Revised: 5 August 2026

Accepted: 14 August 2026

Published: 17 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](#)

[Attribution \(CC BY\)](#) license.

1. Introduction

The transition from federated avionics to Integrated Modular Avionics (IMA) replaced dedicated per-function hardware with shared core processing modules (CPMs), deterministic networks (AFDX per ARINC 664 Part 7 and, increasingly, Time-Sensitive Networking, TSN), and remote data concentrators (RDCs). Watkins and Walter [1] documented this architectural transition, Gaska et al. [2] synthesized its historical evolution and future direction, Annighöfer and Thielecke [3] provided a systems-architecting framework for distributed IMA, and Łukić et al. [4] reviewed the enabling technologies and development priorities. Certification of such systems (DO-297, ARP4761) rests on design-time determinism. Fleming and Leveson [5] addressed hazard analysis and certification; Chen et al. [6] formalized scheduling of independent partitions; Burns and Davis [7] synthesized mixed-criticality scheduling research; and Łukić et al. developed automated ARINC 653 configuration [8] and subsequent validation procedures [9]. Together, these contributions support robust time and space partitioning, worst-case network latency bounds, and design assurance levels (DALs) assigned per function. Yet the very sharing that makes IMA efficient creates coupling channels that certification analyses treat only locally: a queue that saturates on one switch port back-pressures neighboring flows; a degraded partition monopolizes a health-monitoring or reconfiguration resource; a reconfiguration that fails under load leaves faults unconfined. AFDX studies optimized virtual-link design and configuration [10,11], quantified end-to-end response time [12], and evaluated redundant routing, determinism, and reliability [13,14]; complementary TSN studies validated an avionics testbed [15] and implemented a deterministic Ethernet bus [16]. Collectively, these results motivate analysis of emergent, system-level escalation—where a minor local fault propagates into loss of multiple DAL-A functions—that is not excluded by component-level determinism. We refer to the terminal state of such escalation as a catastrophic scenario: the simultaneous loss of the set of highest-criticality functions beyond a certification threshold.

Cascade phenomena on networks have been studied intensively since the load-redistribution model of Motter and Lai [17] and the interdependent-network analysis of Buldyrev et al. [18]. Brummitt et al. [19] proposed a load-cascade suppression strategy; Gao et al. analyzed robustness [20] and structure [21] in networks of networks; and Kivelä et al. [22] and Boccaletti et al. [23] established complementary multilayer-network frameworks. Threshold conditions of the form $R < 1$ are classical in network epidemics [24]. These models, however, are memoryless: the propagation kernel is Markovian, and the state jump rules are instantaneous. Two well-documented features of avionics degradation and traffic violate this assumption. First, functional degradation of computing modules is hereditary: accumulated thermal, radiation-induced, and software-aging effects make the current degradation rate depend on the loading history. Fractional-order operators provide a natural representation of this hereditary behavior: Li et al. established Mittag–Leffler stability [25] and a generalized Lyapunov direct method [26], while Caputo and Fabrizio [27] demonstrated a fractional-derivative model of damage and fatigue. Second, network traffic and retransmission processes exhibit self-similarity and heavy tails [28], for which fractional (non-Markovian) queue dynamics are a parsimonious continuum description [29]. Fractional-order models of networked systems with delays have accordingly become an active area: Zhang et al. [30] derived a Razumikhin-type stability result for delayed memristive neural networks; Zhang et al. [31] treated delayed competitive neural networks; Liu et al. [32] addressed S-type distributed delays; and Cao et al. [33] examined practical stability under distributed delays and impulsive control. Delayed fractional oscillators and their Hopf-type bifurcations have been analyzed [34,35], and qualitative theories of nonlinear delayed fractional systems are now available [36], alongside a growing body of applied fractional-order dynamical models [37,38]. What is missing is a model that couples

this fractional network machinery to the specific escalation mechanisms of IMA—queue accumulation, mixed-criticality priority conflict, and reconfiguration failure—and derives from it certifiable stability margins.

This paper closes that gap. We propose, to our knowledge, for the first time, a graph-based fractional-order model with delays in which four escalation mechanisms of the IMA catastrophic scenario are unified in a single dynamical operator: (M1) network propagation of degradation along information-dependency edges; (M2) queue-backlog accumulation and spillover (backpressure); (M3) blocking of lower-priority partitions by higher-priority ones under resource contention (mixed criticality per ARINC 653); and (M4) state-dependent failure of the reconfiguration mechanism, which we treat as an internal supervisory control loop whose authority collapses under global load. The contributions are:

- A unified fractional-order network model (Section 3): A $2n$ -dimensional system of delayed Caputo fractional differential equations on a weighted digraph of CPM/SW/RDC nodes with DAL-differentiated parameters, in which mechanisms M1–M4 enter as structurally identifiable terms, and whose nominal equilibrium is consistent with nonzero steady-state backlog.
- A cascade threshold number R_c for delayed fractional network dynamics (Theorem 2): $R_c = \rho(M^{-1}(B \circ A + \Gamma))$ aggregates all four mechanisms; $R_c < 1$ implies delay-independent local stability and global Mittag–Leffler stability of the nominal regime, giving a certifiable design margin $1 - R_c$.
- A memory-stabilization result (Theorems 3 and 4): When $R_c > 1$, the sufficient stability certificate is lost and, under the additional conditions stated in Theorem 3, the system is bistable. A critical delay τ^* triggers oscillatory queue precursors of the cascade, and $\tau^*(\alpha)$ is monotonically decreasing in the fractional order—hereditary degradation dynamics *enlarge* the stability region, so the classical integer-order model ($\alpha = 1$) yields the most conservative bound on admissible network delays.
- A synthetic quantitative case study (Section 6): A 22-node IMA configuration simulated with a Grünwald–Letnikov scheme for delayed fractional systems reproduces fault absorption ($R_c = 0.82$), reconfiguration-contained cascades and global catastrophe ($R_c = 1.34$, $T_{cat} \approx 101$ normalized time units), a sharp tipping of the time-to-catastrophe at a critical reconfiguration-capacity threshold ξ^* , and cascade suppression under strong memory ($\alpha = 0.6$)—yielding actionable guidance for reconfiguration-policy design.

The remainder of the paper is organized as follows. Section 2 formalizes the IMA architecture as a graph and the escalation mechanisms. Section 3 presents the fractional-order model. Section 4 presents the main theoretical results together with expanded proofs. Section 5 describes the numerical method. Section 6 reports the simulation study and presents the baseline parameters; Section 7 discusses implications, comparisons, and limitations; and Section 8 concludes.

2. IMA Architecture as a Weighted Digraph and Escalation Mechanisms

2.1. Graph Representation

Let $G = (V, E, W)$ be a weighted directed graph with $|V| = n$ nodes of three types: core processing modules (CPMs) hosting ARINC 653 partitions; AFDX/TSN switches (SW); and remote data concentrators (RDCs) of a distributed IMA. An edge $(j, i) \in E$ with weight $a_{ij} > 0$ encodes an information dependency of node i on node j (a virtual link, a partition-to-partition data flow, or a sensor-feed path); $A = [a_{ij}]$ is the weighted adjacency matrix. Each node carries a criticality class $\text{DAL}(i) \in \{A, B, C\}$ with scheduling priority

$p_i \in \{3, 2, 1\}$, respectively, and a criticality weight w_i (normalized, $\sum_i w_i = 1$) used in system-level load aggregation. For each node i , the contention set

$$H_i = \{k \in V : p_k > p_i, k \text{ and } i \text{ share a CPM or a switch output port}\}$$

collects the higher-priority nodes able to block i under resource contention.

2.2. State Variables and Escalation Mechanisms

Each node i carries two coupled states: the functional degradation level $x_i(t) \in [0, 1]$ (0—nominal, 1—complete loss of the hosted function) and the normalized queue backlog $q_i(t) \geq 0$ (aggregated occupancy of the node's transmit/receive buffers relative to a design reference). Four mechanisms couple these states across the graph:

M1 (Network Propagation). Degradation of an upstream node j corrupts, delays, or invalidates the data consumed by i , raising the effective degradation of i after the communication delay τ_{ij} .

M2 (Backlog Spillover). Saturated buffers at j back-pressure the shared switch fabric, inflating the backlog at neighboring i after delay σ_{ij} ; simultaneously, degradation of i reduces its own service capacity, so backlog and degradation reinforce each other.

M3 (Priority Conflict). Under contention, partitions with $p_k > p_i$ preempt CPU windows or output-port bandwidth; the blocking experienced by i grows with both its own backlog and the degradation of the blockers (a degraded high-priority partition holds resources longer).

M4 (Reconfiguration Failure). IMA health monitoring migrates or restarts degraded functions at rate r_i ; this supervisory loop itself consumes computing and network resources, and its success probability collapses when the criticality-weighted global load exceeds a capacity threshold.

A catastrophic scenario is the event that all DAL-A functions degrade beyond a certification threshold: with $C_A = \{i : \text{DAL}(i) = A\}$ and $x_{cat} \in (0, 1)$,

$$\mathcal{K} = \{x \in [0, 1]^n : \min_{i \in C_A} x_i \geq x_{cat}\}, \quad T_{cat} = \inf\{t \geq 0 : x(t) \in \mathcal{K}\}. \quad (1)$$

Architectural Origin of the Coupling Terms

The four coupling terms are reduced-order representations of identifiable mechanisms in the underlying IMA architecture. The degradation-propagation term follows the directed information-dependency graph: an edge $j \rightarrow i$ represents a virtual link, partition dependency, or sensor-data path through which degradation of node j may impair the function hosted by node i . The corresponding weight quantifies the strength of this dependency and may be derived from data criticality, traffic share, redundancy, and functional dependence.

The queue-backlog term represents the balance between traffic arrival and effective service capacity. Degradation reduces processing or transmission capability, while congestion at adjacent nodes produces spillover or backpressure through shared switching and routing resources. The priority-conflict term represents mixed-criticality contention under ARINC 653 scheduling, where higher-priority or degraded partitions may occupy processor windows or communication resources for longer intervals and thereby delay lower-priority functions.

The reconfiguration term represents the Health Monitoring and recovery subsystem as a finite-capacity supervisory loop. Its nominal rate reflects the available restart, migration, or isolation capability, whereas its load-dependent reduction captures the loss of effectiveness caused by processor saturation, queue congestion, or competition for recovery resources.

Accordingly, the graph determines where interactions are possible, while the coupling coefficients determine their intensity. These coefficients should be calibrated from architecture descriptions, scheduling and network analyses, telemetry, and reconfiguration logs. The present formulation is therefore a phenomenological system-level reduction that preserves the dominant IMA mechanisms while remaining analytically tractable.

3. Fractional-Order Model

3.1. Preliminaries

For $0 < \alpha \leq 1$, the Caputo derivative of $y \in C^1$ is

$${}^C D^\alpha y(t) = \frac{1}{\Gamma(1-\alpha)} \int_0^t (t-s)^{-\alpha} y'(s) ds, \quad (2)$$

reducing to y' at $\alpha = 1$. The two-parameter Mittag–Leffler function is $E_{\alpha,\beta}(z) = \sum_{k=0}^{\infty} z^k / \Gamma(\alpha k + \beta)$, $E_\alpha := E_{\alpha,1}$ [25,26]. Matignon’s criterion states that the linear system ${}^C D^\alpha y = Jy$ is asymptotically stable iff $|\arg \lambda_i(J)| > \alpha\pi/2$ for all eigenvalues [39]; the sector widens as α decreases, which is the linear seed of the memory-stabilization effect quantified in Theorem 4.

The choice of Caputo operators of orders α (degradation) and β (backlog) is a modeling decision with empirical grounding: (i) degradation of avionics computing modules is hereditary—aging, thermal, and single-event-effect accumulation make the instantaneous degradation rate a power-law-weighted functional of the loading history, the defining property of the kernel in Equation (2) [26,27]; and (ii) LAN and avionics-bus traffic exhibits self-similarity and heavy-tailed burst/retransmission statistics [28], for which fractional relaxation of the backlog is the natural continuum limit [29]. Both orders are treated as identifiable parameters to be calibrated from telemetry; the integer case $\alpha = \beta = 1$ is retained throughout as the comparison baseline.

3.1.1. Advantages over the Integer-Order Counterpart

The fractional-order formulation is a nested extension of the corresponding integer-order model rather than an unrelated alternative. Setting $\alpha = \beta = 1$ recovers the conventional integer-order dynamics while preserving the graph topology, communication delays, coupling coefficients, nonlinearities, and reconfiguration mechanism. This property enables a controlled comparison in which the influence of hereditary memory can be separated from changes in the underlying IMA architecture.

For $0 < \alpha < 1$, the Caputo derivative of the degradation state can be written as

$${}^C D_t^\alpha x_i(t) = \frac{1}{\Gamma(1-\alpha)} \int_0^t (t-s)^{-\alpha} \dot{x}_i(s) ds. \quad (3)$$

The power-law kernel shows that the current degradation dynamics depend on the complete loading and fault history rather than only on the instantaneous state. Consequently, two IMA nodes with the same current degradation level may evolve differently if they have experienced different histories of thermal stress, queue congestion, transient faults, or incomplete recovery. A conventional integer-order model of the same state dimension does not reproduce this hereditary effect unless additional latent states, distributed delays, or a high-order approximation are introduced.

Fractional-order dynamics also permit non-exponential Mittag–Leffler relaxation, which may represent long-tailed recovery and persistent degradation more compactly than a finite-dimensional integer-order model. The fractional orders therefore act as identifiable memory parameters: values close to one represent weak hereditary behavior, whereas lower values represent stronger persistence of previous system states. This additional

modeling flexibility is useful only when supported by data; if the estimated orders are statistically indistinguishable from one, the integer-order formulation remains the appropriate reduced model.

The fractional-order formulation is not assumed to be necessary a priori; it is treated as a testable reduced-order hypothesis for hereditary IMA dynamics. The integer-order counterpart is recovered by setting the derivative orders to one. For fractional orders below one, the Caputo kernels represent the long-tailed influence of previous loading, faults, congestion, and recovery. In the degradation channel, this memory may arise from accumulated thermal stress, software aging, repeated transient disturbances, and incomplete recovery. In the backlog channel, it may represent persistent traffic bursts, retransmissions, correlated arrivals, and non-exponential queue relaxation.

This formulation allows two states with identical instantaneous values to evolve differently because of their different histories, which a same-dimensional Markovian integer-order model cannot reproduce without additional latent states or distributed delays. It also permits Mittag–Leffler relaxation and order-dependent delay-stability margins. Nevertheless, the fractional model should be retained only if telemetry demonstrates a reproducible predictive advantage over the corresponding integer-order baseline; otherwise, the reduced integer-order case is preferable.

3.1.2. Empirical Status and Identification of the Fractional Orders

The fractional-order formulation should be interpreted as a physically motivated modeling hypothesis rather than as an already validated representation of a particular certified IMA platform. The present study does not use proprietary aircraft telemetry or manufacturer-specific measurements from Airbus, Boeing, or another avionics supplier. Publicly accessible datasets containing synchronized, sufficiently long records of component degradation, queue occupancy, network traffic, and reconfiguration outcomes are limited; therefore, the fractional orders used in the synthetic case study should not be interpreted as universal constants or as values established for operational aircraft.

The motivation for introducing fractional-order dynamics is based on the possibility of hereditary behavior in both state channels. For the degradation state, repeated thermal loading, accumulated component stress, software aging, and recovery after transient faults may cause the current degradation rate to depend on the previous workload and fault history. For the backlog state, persistent traffic bursts, retransmission sequences, and correlated workload variations may produce relaxation and dependence structures that differ from the exponential behavior assumed by an integer-order Markovian model. A Caputo derivative represents these effects through a power-law memory kernel and therefore provides a parsimonious alternative to a high-dimensional integer-order model containing numerous latent memory states.

This physical motivation alone is insufficient to establish that fractional-order dynamics are necessary for every IMA implementation. The orders α and β must therefore be treated as identifiable architecture- and operating-condition-dependent parameters. Their empirical estimation would require representative health-monitoring and network-telemetry records obtained under nominal operation, controlled workload changes, and fault-injection or degradation experiments. For the degradation channel, an integer-order exponential relaxation model and a fractional Mittag–Leffler relaxation model can be fitted to the measured post-disturbance response. For the backlog channel, the analysis can additionally examine autocorrelation decay, burst persistence, and scaling behavior over multiple time intervals.

The fractional formulation should be retained only when it provides a reproducible improvement over the corresponding integer-order model, for example, in terms of out-of-

sample prediction error, residual correlation, or information criteria. If the estimated orders are statistically indistinguishable from one another, the integer-order model constitutes the appropriate reduced representation. Accordingly, the present simulations demonstrate the qualitative consequences of memory-dependent dynamics and provide a calibration framework, but they do not constitute manufacturer-specific empirical validation.

3.1.3. Interpretation and Dependence of the Fractional Orders

Although the degradation state x_i and the queue-backlog state q_i are governed by equations with a similar graph-coupled structure, they represent physically distinct processes and are therefore not required to have the same fractional order. The order α characterizes hereditary effects in functional degradation, such as accumulated thermal or computational stress, software aging, repeated transient disturbances, and non-exponential recovery. The order β characterizes temporal persistence in queue evolution, including correlated traffic bursts, retransmission sequences, long-range dependence, and non-Markovian backlog relaxation. The common mathematical structure reflects the shared IMA dependency graph and the bidirectional coupling of the two state channels rather than an assumption of identical memory mechanisms.

In the present model, α and β are treated as separately identifiable parameters. This assumption does not imply independence of the degradation and backlog dynamics. Their interaction is represented explicitly in the governing equations: degradation reduces the effective service capability and promotes queue growth, whereas excess backlog increases the degradation rate and contributes to cascade escalation. Consequently, the two channels are dynamically cross-coupled even though their Caputo kernels are specified separately.

The current formulation assumes separable power-law memory kernels and does not include direct cross-memory terms in which the historical trajectory of q_i modifies the fractional kernel of x_i , or vice versa. This modeling choice preserves analytical tractability and allows the two orders to be estimated from different classes of telemetry. A more general formulation could employ coupled, matrix-valued, distributed-order, or variable-order operators; however, such models would introduce additional parameters that cannot be reliably identified without synchronized, long-duration degradation and network-traffic measurements.

For platform-specific calibration, α and β should first be estimated separately by fitting the degradation and queue responses to representative telemetry. Their joint uncertainty and statistical dependence should then be assessed. If both estimates are indistinguishable within their confidence intervals, the reduced assumption $\alpha = \beta$ may be adopted. If they differ significantly, retaining distinct orders provides a more faithful representation of the two memory mechanisms. A systematic correlation between the estimated orders and a common operating variable would indicate the need for a constrained or coupled-memory extension.

Thus, fractional-order dynamics are used here as a testable memory hypothesis, not as an assumed universal property of IMA systems.

3.2. Governing Equations

For $i = 1, \dots, n$, with initial functions $\varphi_i \in C([- \tau_{\max}, 0])$:

$$\begin{aligned} {}^C D^\alpha x_i(t) &= -\mu_i x_i(t) + \sum_{j=1}^n \beta_{ij} a_{ij} f(x_j(t - \tau_{ij})) + \gamma_i g((q_i(t) - \bar{q}_i)_+) + \eta_i \Phi_i(t) - \rho_i(t) x_i(t), \\ {}^C D^\beta q_i(t) &= \lambda_i - c_i(1 - x_i(t)) h(q_i(t)) + \sum_{j=1}^n w_{ij} q_j(t - \sigma_{ij}) - \delta_i q_i(t), \end{aligned} \quad (4)$$

where $\alpha, \beta \in (0, 1]$; $\mu_i > 0$ is the intrinsic recovery rate; $\beta_{ij} \geq 0$ the degradation-transmission intensity along edge (j, i) ; $\tau_{ij}, \sigma_{ij} \in [0, \tau_{\max}]$ communication delays; $\lambda_i > 0$ the

traffic arrival intensity; $c_i(1 - x_i)$ the service capacity, which collapses with degradation; $w_{ij} \geq 0$ the backlog-spillover (backpressure) coefficients; $\delta_i > 0$ the backlog drain; and $(u)_+ = \max(u, 0)$. The saturating nonlinearities f, g, h are smooth, nondecreasing, and Lipschitz with $f(0) = g(0) = h(0) = 0$, $f'(0) = g'(0) = 1$; we use $f(u) = g(u) = u/(1 + u)$ and $h(u) = u/(1 + u/2)$. The reference backlog $\bar{q}_i$ is the nominal steady-state occupancy: only excess backlog degrades functionality, which makes the nominal pair $E_0 = (0, q^*)$ an exact equilibrium (Section 4).

Priority-conflict term (M3). Blocking of node i by its contention set:

$$\Phi_i(t) = q_i(t) \sum_{k \in H_i} \theta_{ik} \frac{x_k(t - \tau_{ik})}{1 + x_k(t - \tau_{ik})}, \quad (5)$$

so blocking requires both pending load at i ($q_i > 0$) and degraded blockers ($x_k > 0$), saturating in x_k . Nodes of the highest priority have $H_i = \emptyset$ and experience no blocking, consistent with ARINC 653 scheduling.

Reconfiguration failure (M4). The reconfiguration (suppression) rate is $\rho_i(t) = r_i(1 - \pi(t))$, where r_i is the DAL-differentiated health-monitoring capacity (r_i largest for DAL-A partitions, which are restored first) and

$$\pi(t) = (1 + \exp[-\kappa(\langle x(t) \rangle_w + \nu \langle q(t) \rangle_w - \xi)])^{-1}, \quad \langle y \rangle_w = \sum_i w_i y_i, \quad (6)$$

is the probability that reconfiguration fails, rising sigmoidally once the criticality-weighted global load exceeds the capacity threshold ξ with steepness κ . Equation (6) is a deliberately aggregated (phenomenological) closure: it is the mean-field limit of a two-state Markov availability model of the reconfiguration mechanism whose repair rate is load-dependent, and we position it as such—Section 7 discusses this reduction and its calibration. Structurally, $\rho_i(t)x_i$ is an internal state-feedback suppression loop: Equation (6) is the point where supervisory control authority enters the dynamics, and its collapse ($\pi \rightarrow 1$) is what converts a contained fault into a catastrophic scenario.

3.2.1. Interpretation and Possible Distributed Extension of the Health Monitoring Model

The sigmoid function in Equation (6) is an aggregated representation of Health Monitoring and reconfiguration availability rather than a detailed model of the complete distributed HM architecture. Its purpose is to capture a system-level loss of supervisory effectiveness when the criticality-weighted load approaches the finite capacity of the monitoring and reconfiguration resources. Accordingly, the parameters of the sigmoid characterize the location and sharpness of an aggregate capacity transition and should not be interpreted as direct representations of individual HM components or software processes.

Although Equation (6) depends explicitly on the global weighted load, the latter is generated by spatially heterogeneous dynamics on the IMA graph. Local degradation, queue accumulation, communication delays, routing dependencies, service-capacity reduction, and mixed-criticality contention determine the contribution of each node to the aggregate load. The HM closure therefore receives a topology-dependent system state, even though the resulting failure probability is represented by a scalar mean-field function. DAL-differentiated reconfiguration rates additionally represent the preferential restoration of higher-criticality functions.

This aggregation is appropriate for the present objective of deriving a tractable system-level cascade threshold and identifying the tipping of reconfiguration capacity. However, it does not explicitly distinguish among local HM processors, recovery queues, routing paths, redundant monitoring channels, or node-specific resource constraints. For a platform-

specific implementation, Equation (6) may be replaced by a distributed reconfiguration model of the form

$$r_i^{\text{eff}}(t) = r_i^0 A_i^{\text{HM}}(t) [1 - P_i^{\text{fail}}(t)], \quad (7)$$

where $A_i^{\text{HM}}(t)$ denotes the availability of the HM and reconfiguration path serving node i , and $P_i^{\text{fail}}(t)$ may depend on the local backlog, available recovery capacity, routing reachability, and the states of the supervisory nodes:

$$P_i^{\text{fail}}(t) = \sigma \left(\kappa_i \left[w_i^{\text{loc}} q_i(t) + \sum_{j \in \mathcal{H}_i} w_{ij}^{\text{HM}} x_j(t) + w_i^{\text{res}} u_i(t) - \Theta_i \right] \right). \quad (8)$$

Here, $\mathcal{H}_i$ denotes the set of HM, communication, or computational resources required to detect, isolate, and recover node i ; $u_i(t)$ is the utilization of the corresponding recovery resources; and Θ_i is a node- or subsystem-specific capacity threshold. This formulation would allow topology, routing, local workload, and resource allocation to affect reconfiguration success explicitly.

The global sigmoid used in the present study can be interpreted as a mean-field reduction in this distributed representation under the assumption that local HM capacities and resource states are sufficiently homogeneous or can be summarized by the criticality-weighted global load. This assumption enables analytical treatment but must be validated before applying the model to a specific certified IMA platform.

3.2.2. Physical Interpretation of the Sigmoid Steepness Parameter

The parameter κ controls the steepness of the transition between effective and ineffective reconfiguration as the criticality-weighted system load approaches the available Health Monitoring capacity. It should therefore be interpreted as a measure of the sensitivity of the supervisory infrastructure to resource saturation rather than as a physical property of an individual hardware component.

Small values of κ describe systems in which reconfiguration effectiveness decreases gradually with increasing load. Conversely, large values correspond to architectures that remain almost fully operational until the available computational or communication resources become nearly exhausted, after which the probability of successful recovery decreases rapidly. Such behavior is consistent with many resource-limited distributed systems, where scheduler overload, queue saturation, processor utilization, or communication congestion may produce a relatively abrupt degradation of service once a critical utilization threshold is exceeded.

The value $\kappa = 14$ adopted in the present synthetic case study was selected as a representative parameter that produces a sufficiently sharp but still continuously differentiable transition. Although the resulting sigmoid closely approximates threshold behavior near the capacity limit, it remains smooth and therefore preserves the assumptions required by the theoretical analysis presented in Section 4. The parameter should not be interpreted as a universal engineering constant and is expected to depend on the implementation of the Health Monitoring and reconfiguration subsystem.

For platform-specific applications, κ should be identified empirically from Health Monitoring records by analyzing the relationship between the probability of successful reconfiguration and the corresponding criticality-weighted system load. The calibration procedure may employ logistic regression or other nonlinear parameter-estimation techniques using observations collected over a broad range of operating conditions.

The full model, Equations (3)–(6), is a $2n$ -dimensional system of nonlinear delayed fractional differential equations on $\Omega = [0, 1]^n \times \mathbb{R}_+^n$.

4. Main Theoretical Results

4.1. Standing Assumptions and Scope of the Results

Throughout this section, the graph is finite and weighted by a nonnegative adjacency matrix; all delays are fixed and bounded; the initial histories are continuous and belong to the feasible state domain; and all recovery, service, and reconfiguration parameters satisfy the positivity and inward-pointing conditions required for invariance. The nonlinear propagation, backlog, priority-conflict, and reconfiguration functions are continuously differentiable, bounded on the feasible domain, and globally Lipschitz there. Additional equilibrium, hyperbolicity, simple-root, and transversality assumptions are stated separately for Theorems 3 and 4.

The cascade-threshold condition is a sufficient delay-independent stability certificate. Conversely, loss of this certificate does not, by itself, imply bistability. The bistability result is therefore conditional on the additional assumptions of Theorem 3. The memory-stabilization theorem is proved for the scalar dominant-mode reduction; the corresponding network-level effect is evaluated numerically.

For clarity and mathematical traceability, each theorem is followed immediately by an expanded proof that states the principal logical steps and conditions. All equations required by these arguments are displayed and numbered in this section.

Theorem 1 (Well-posedness and Positive Invariance). *For any initial function $\varphi \in C([- \tau_{\max}, 0], \Omega)$, the system of Equations (3)–(6) admits a unique solution on $[0, \infty)$, and Ω is positively invariant: $x_i(t) \in [0, 1]$ and $q_i(t) \geq 0$ for all $t \geq 0$, provided $\lambda_i > 0$ and the vector field satisfies the stated Lipschitz conditions.*

$$z = (x, q) \in \Omega \quad {}^C D^{\bar{\alpha}} z(t) = F(z(t), z(t - \tau.)) \bar{\alpha} \in \{\alpha, \beta\} f, g, h \Omega \Phi_i \pi F \Omega L_F \quad (9)$$

Proof.

$$z_i(t) = \varphi_i(0) + \frac{1}{\Gamma(\bar{\alpha}_i)} \int_0^t (t-s)^{\bar{\alpha}_i-1} F_i(z(s), z(s-\tau.)) ds.$$

On with the weighted norm, the right-hand side of the integral equation above is a contraction for large (the delay terms only refer to already-determined history on, handled by the method of steps), yielding a unique local solution; a priori boundedness via the generalized Gronwall inequality [40] gives global continuation. For invariance: on the face, all terms of Equation (3) except (which vanish) are, so the fractional derivative of is and the vector field points inward; on, using, and the normalization of Table 1, the field satisfies. The invariance criterion for Caputo systems (fractional Nagumo condition; see [36]) completes the proof. $\square$

Theorem 2 (Cascade Threshold and Global Mittag–Leffler Stability). *Let $E_0 = (0, q^*)$ be the nominal equilibrium, where q^* solves $\lambda_i = c_i h(q_i^*) + \delta_i q_i^* - \sum_j w_{ij} q_j^*$. Define*

$$R_c = \rho \left(M^{-1} (B \circ A + \Gamma) \right), \quad M = \text{diag}(\mu_i + r_i(1 - \pi_0)),$$

where $B \circ A = [\beta_{ij} a_{ij}]$ and Γ collects the linearized sensitivities of the q - and Φ -channels at E_0 , as specified in the proof below. If $R_c < 1$, then E_0 is locally asymptotically stable for arbitrary delays $\tau_{ij}, \sigma_{ij} \geq 0$ (delay-independent stability), and globally Mittag–Leffler stable in Ω :

$$\|x(t)\| \leq C \|\varphi\| E_{\alpha}(-\omega t^{\alpha}), \quad \omega > 0.$$

$$E_0 = (0, q^*) f'(0) = g'(0) = 1x$$

Proof.

$${}^C D^\alpha \tilde{x} = -M\tilde{x} + (B \circ A) \tilde{x}(t - \tau) + \Gamma_q \tilde{q} + \Gamma_\Phi \tilde{x}(t - \tau), \quad \Gamma := \Gamma_q S + \Gamma_\Phi,$$

where $\Gamma_\Phi = [\eta_i q_i^* \theta_{ik} \mathbb{1}_{k \in H_i}]$ is the priority-conflict sensitivity, $\Gamma_q = \text{diag}(\gamma_i)$, and S is the steady-state sensitivity of the (input-to-state stable) q -subsystem to x at E_0 , bounded elementwise by the q -block Jacobian inverse. Suppose $R_c < 1$. Then, $M - (B \circ A + \Gamma)$ is a nonsingular M-matrix, so $v \gg 0$ with $v^\top (M - (B \circ A + \Gamma)) > 0$ exists componentwise. Define $V(x) = v^\top x \geq 0$ on $[0, 1]^n$. By the Caputo differentiation lemma for linear functionals and the concavity of f, g ($f(u) \leq u, g(u) \leq u$ on $\mathbb{R}_+$), along solutions

$${}^C D^\alpha V \leq -v^\top Mx(t) + v^\top (B \circ A + \Gamma) \bar{x}_t, \quad \bar{x}_t := \sup_{s \in [t - \tau_{\max}, t]} x(s) \text{ (componentwise)}.$$

Under the Razumikhin condition, for some close to (possible since leaves slack), the preceding differential inequality gives with. The fractional Razumikhin theorem and the comparison principle [30,41] then yield, hence Equation (8). Delay-independence is inherent: no delay bound was used. Global validity follows because the concavity bounds hold on all of and the subsystem is input-to-state stable with respect to by the M-matrix structure of its own linear part (Table 1). $\square$

The quantity $1 - R_c$ is thus a certifiable stability margin: all four escalation mechanisms enter Equation (7) through $B \circ A$ (M1), the q -sensitivities in Γ (M2), the Φ -sensitivities in Γ (M3), and the reconfiguration authority $r_i(1 - \pi_0)$ in M (M4).

Table 1. Baseline parameters of the synthetic 22-node IMA case study. Time is dimensionless and normalized by one Health Monitoring major frame; physical-delay equivalents are provided where applicable.

Symbol	Meaning	Value
n	nodes (CPM/SW/RDC)	22 (8/4/10)
DAL A/B/C	criticality classes	4/8/10 nodes
α, β	fractional orders	$\{0.6, 0.7, 0.8, 0.9, 1.0\}$, baseline 0.8
μ_i	intrinsic recovery	0.6
β_{ij}	transmission intensity	$\varkappa \cdot 0.25 a_{ij}$; $\varkappa = 0.93$ (S1), 1.60 (S2/S3)
γ_i	backlog $\rightarrow$ degradation gain	0.15
η_i	priority-conflict gain	0.22
θ_{ik}	conflict intensity	0.4
r_i	reconfiguration rate	DAL A/B/C: 1.6/0.6/0.25
λ_i	traffic arrival	$U(0.25, 0.35)$
c_i	service capacity	1.1
δ_i	backlog drain	0.3
w_{ij}	backlog spillover	$0.04 \mathbb{1}_{a_{ij} > 0}$
κ, ν, ξ	Equation (6)	14, 0.35, 0.38–0.75 per scenario
x_{cat}	certification threshold	0.6
τ_{ij}, σ_{ij}	delays	$[0.3, 4]$ (physical 0.5–3 ms)
h, T	grid step, horizon	0.05, 200–400

4.2. Interpretation and Conservatism of the Cascade Threshold

The condition $R_c < 1$ is a sufficient stability condition and should not be interpreted as an exact nonlinear bifurcation boundary. Its satisfaction guarantees global Mittag–Leffler stability of the nominal equilibrium under the assumptions of Theorem 2. By contrast, $R_c \geq 1$ indicates that the M-matrix and Lyapunov–Razumikhin certificate are no longer available; it does not imply that every initial history necessarily converges to a catastrophic

state. The threshold therefore separates a guaranteed stable region from a potentially bistable or cascade-prone region.

The condition is deliberately conservative because the nonlinear system is bounded by a positive comparison model. In particular, Γ contains upper bounds on the sensitivities of the backlog and priority-conflict channels, and the proof employs global monotonicity and Lipschitz bounds over the feasible domain. Moreover, the delay-independent argument does not use the actual magnitudes or distributions of τ_{ij} and σ_{ij} . These approximations provide a robust one-sided certificate but may underestimate the true stability region of a particular architecture.

This interpretation is consistent with the numerical scenarios. For $R_c = 0.82$, Scenario S1 remains within the analytically certified region and the initiating disturbance is absorbed. However, $R_c = 1.34$ is used in both Scenarios S2 and S3: the disturbance is contained in S2 while the reconfiguration mechanism remains effective, whereas S3 reaches the catastrophic set after the supervisory capacity collapses. Thus, $R_c > 1$ identifies the possibility of cascade escalation but does not constitute a deterministic prediction of catastrophe.

The conservatism of the threshold is architecture- and operating-point-dependent. For a calibrated platform, it may be quantified by comparing the analytical boundary $R_c = 1$ with a transition boundary identified through numerical continuation, characteristic-root analysis, ensemble simulation, or representative fault-injection experiments.

4.3. Structural Threshold and Delay–Memory Stability Boundary

Two distinct but complementary stability boundaries should be distinguished. The cascade threshold R_c in Equation (7) is a structural amplification criterion derived from a positive comparison system. It is deliberately independent of the communication delays and fractional orders because the Lyapunov–Razumikhin proof bounds the delayed states over the admissible history interval. Therefore, $R_c < 1$ provides a sufficient Mittag–Leffler stability certificate for arbitrary fixed delays covered by the assumptions of Theorem 2. Within this region, delay and fractional memory affect the transient trajectory and convergence rate but do not change the guaranteed stability classification.

When $R_c \geq 1$, the sufficient comparison certificate is lost, and stability must be evaluated from the characteristic roots of the delayed fractional-order system. For a homogeneous delay, the linearized characteristic equation has the generic form

$$\Delta(s; \alpha, \beta, \tau) = \det \left[\begin{pmatrix} s^\alpha I_n & 0 \\ 0 & s^\beta I_n \end{pmatrix} - J_0 - J_\tau e^{-s\tau} \right] = 0,$$

where J_0 and J_τ denote the instantaneous and delayed Jacobian components, respectively. At a stability crossing $s = i\omega$,

$$(i\omega)^\nu = \omega^\nu \left[\cos\left(\frac{\nu\pi}{2}\right) + i\sin\left(\frac{\nu\pi}{2}\right) \right], \quad \nu \in \{\alpha, \beta\},$$

while the delayed coupling contributes the phase factor

$$e^{-i\omega\tau} = \cos(\omega\tau) - i\sin(\omega\tau).$$

The critical delay is therefore determined by a joint magnitude and phase balance involving the fractional orders and the delayed network feedback. Increasing τ increases the phase lag of the delayed coupling, whereas decreasing a fractional order changes both the dynamic magnitude and the fractional stability-sector geometry.

Accordingly, R_c and τ^* have different interpretations. The condition $R_c < 1$ defines a delay-independent guaranteed stability region, whereas $\tau^* = \tau^*(\alpha, \beta)$ defines a delay–

memory-dependent local stability boundary when the structural certificate is unavailable. Crossing $R_c = 1$ does not itself imply an immediate catastrophe; rather, it makes the system susceptible to the delay-dependent and initial-condition-dependent behavior characterized in Theorems 3 and 4.

Theorem 3 (Conditional Bistability and Delay-Induced Oscillatory Instability). *Suppose the sufficient stability certificate is unavailable and, in addition, the equilibrium-existence, local-stability, hyperbolicity, basin-separation, simple-root, and transversality assumptions stated above hold. Then, the system admits, in addition to at least one interior catastrophic equilibrium, and the dynamics are bistable. For a homogeneous delay, the characteristic function of the linearization at the relevant equilibrium, $R_c > 1 E_0 E^* \min_{i \in C_A} x_i^* > 0 \tau_{ij} = \sigma_{ij} = \tau$*

$$\Delta(s) = \det (s^\alpha I - J_0 - J_\tau e^{-s\tau}),$$

admits a critical delay $\tau^* > 0$ at which a root pair crosses the stability sector $|\arg s| = \alpha\pi/2$ transversally (fractional Hopf-type bifurcation), producing sustained queue oscillations that precede cascade onset.

$$\begin{aligned} R_c > 1 T : [0, 1]^n \rightarrow [0, 1]^n x T(x)_i x T \Omega \geq 0 T(0) = 0 D T(0) R_c > 10 E^* \gg 0 T(1) \leq 1 E^* \\ \in (0, 1]^n E_0 E_0 s = \omega e^{i\pi/2} \omega > 0 a b J_0 J_\tau d \operatorname{Re} s / d \tau|_{\tau^*} > 0 \Delta(s(\tau), \tau) = 0 \partial_s \Delta \\ \neq 0 \square \end{aligned}$$

Proof. The statement is conditional on the additional equilibrium-existence, hyperbolicity, and basin-separation assumptions stated above. Let E_0 denote the nominal equilibrium and let E^* denote an interior catastrophic equilibrium satisfying $x_i^* \geq x_{CAT}$ for all $i \in C_A$, where C_A is the set of DAL-A functions. The existence of E^* follows from the assumed existence of an interior solution of the equilibrium equations associated with Equations (3)–(6) in which the reconfiguration effectiveness is reduced by the load-dependent term in Equation (6). Under this assumption, both E_0 and E^* are equilibria of the same delayed fractional-order graph system.

The local stability of E_0 in the subcritical neighborhood follows from the comparison and Mittag–Leffler stability argument of Theorem 2 whenever the corresponding local comparison matrix satisfies the sufficient threshold condition. For the catastrophic equilibrium E^* , the assumed local stability means that the spectrum of the linearized delayed fractional-order system around E^* lies inside the admissible fractional stability sector. Hence, there exist neighborhoods U_0 and U^* such that trajectories initialized in U_0 converge to E_0 , whereas trajectories initialized in U^* converge to E^* . Since $E_0 \neq E^*$ and both equilibria are locally asymptotically stable under the stated assumptions, their basins of attraction are disjoint open sets in the feasible state domain. The boundary separating these basins defines a separatrix. This proves the conditional bistability claim.

It remains to justify the delay-induced oscillatory instability. Linearizing Equations (3)–(6) around the equilibrium under consideration gives a delayed fractional-order variational system of the form

$${}^c D^\nu z(t) = J_0 z(t) + J_\tau z(t - \tau),$$

where $z(t)$ collects perturbations of the degradation and backlog states, J_0 is the instantaneous Jacobian, J_τ is the delayed-coupling Jacobian, and ν denotes the relevant fractional order or block-diagonal set of fractional orders. For a homogeneous delay, the associated characteristic equation is

$$\Delta(s; \tau) = \det [\operatorname{diag}(s^\alpha I_n, s^\beta I_n) - J_0 - J_\tau e^{(-s\tau)}] = 0.$$

By the stated simple-root and transversality assumptions, there exists a critical delay τ^* and a frequency $\omega^* > 0$ such that

$$\Delta(i\omega^*; \tau^*) = 0,$$

and the corresponding conjugate root pair crosses the fractional stability boundary with nonzero crossing speed,

$$d \operatorname{Re} s(\tau) / d\tau |_{\tau = \tau^*} \neq 0.$$

For $0 < \tau < \tau^*$, the relevant roots remain in the admissible fractional stability sector; when τ crosses τ^* , the transversality condition implies loss of local stability through an oscillatory root crossing. This is the fractional analog of a Hopf-type instability for delayed systems. Because the backlog variables are directly affected by delayed spillover and service-capacity degradation, the unstable oscillatory mode appears first as a queue-backlog oscillation. These oscillations can move the system trajectory across the basin boundary separating E_0 and E^* , thereby acting as precursors of cascade escalation. The result is conditional: loss of the sufficient stability certificate alone does not imply bistability. Bistability and delay-induced oscillatory instability require the additional equilibrium-existence, local-stability, hyperbolicity, simple-root, transversality, and basin-separation assumptions stated in the theorem. This completes the proof. $\square$

Theorem 4 (Memory Stabilization). *For the scalar test system ${}^C D^\alpha y = ay + by(t - \tau)$ with $a < 0 < |b|$, $|b| > |a|$, the critical delay $\tau^*(\alpha)$ obtained from Equation (9) is monotonically decreasing on $\alpha \in (0, 1]$. Consequently, decreasing the fractional order—i.e., strengthening the hereditary memory of the degradation process—enlarges the delay-stability region, and the classical integer-order model $\alpha = 1$ yields the most conservative (smallest) bound on admissible network delays.*

$${}^C D^\alpha y = ay + by(t - \tau)s = i\omega\omega > 0s^\alpha - a - be^{-s\tau} = 0s^\alpha = \omega^\alpha e^{i\alpha\pi/2}$$

Proof.

$$\left| \omega^\alpha e^{i\alpha\pi/2} - a \right| = |b|$$

$$\tau^*(\alpha) = \frac{1}{\omega} \left(\arg \left[\omega^\alpha e^{i\alpha\pi/2} - a \right] - \arg b \right) \pmod{2\pi/\omega}.$$

Write $P(\omega, \alpha) = \omega^{2\alpha} - 2a\omega^\alpha \cos(\alpha\pi/2) + a^2 - b^2$; the crossing condition above reads $P = 0$. Since $a < 0$, $\partial_\omega P = 2\alpha\omega^{\alpha-1}(\omega^\alpha - a\cos(\alpha\pi/2)) > 0$ and $\partial_\alpha P = 2\omega^\alpha \ln\omega(\omega^\alpha - a\cos(\alpha\pi/2)) + a\pi\omega^\alpha \sin(\alpha\pi/2)$. Implicit differentiation gives $d\omega/d\alpha = -\partial_\alpha P / \partial_\omega P$. Substituting into the total derivative of the critical-delay expression above and using $a < 0$ (so the phase $\arg[\omega^\alpha e^{i\alpha\pi/2} - a] \in (0, \pi/2)$ increases in both ω^α and α) while the prefactor $1/\omega$ decreases along the crossing branch as α grows (the crossing frequency $\omega(\alpha)$ increases because the sector $|\arg s| = \alpha\pi/2$ opens toward the imaginary axis), one obtains $d\tau^*/d\alpha < 0$ on the principal branch for all $a < 0 < |b|$, $|b| > |a|$. Numerically, for $a = -0.6$, $b = 0.75$: $\tau^*(0.6) = 3.00$, $\tau^*(0.8) = 1.76$, and $\tau^*(1.0) = 1.43$. $\square$

5. Numerical Method

Delayed fractional systems require dedicated discretizations because the Caputo operator carries the full solution history. We use an explicit Grünwald–Letnikov (GL) scheme: on the uniform grid $t_k = kh$,

$$y_{k+1} = h^\alpha F(y_k, y_{k-d}) - \sum_{j=1}^{k+1} c_j^{(\alpha)} y_{k+1-j}, \quad c_0^{(\alpha)} = 1, \quad c_j^{(\alpha)} = (1 - \alpha + 1/j)c_{j-1}^{(\alpha)},$$

with delays $d = \text{round}(\tau/h)$, step $h = 0.05$, horizon $T \in [200, 400]$, $hL_\infty 10^{-3}$ (deviation $< 10^{-3}$ for physical delays $0.5 - 3$ ms, $\tau \in [0.3, 4]$), with delays realized as integer lags on the grid, and the history sum evaluated by vectorized inner products with an optional short-memory window [42]. The step (dimensionless time) and horizon were verified by (i) halving (relative deviation in all reported scenarios) and (ii) cross-checking scenario S1 against the Adams–Bashforth–Moulton predictor–corrector of Diethelm, Ford, and Freed [43] adapted to fixed delays (deviation). Physical link delays are mapped to dimensionless values by the degradation time-scale normalization stated in the caption of Table 1.

All figures were prepared and verified at their intended publication dimensions, with enlarged axis labels, tick labels, legends, and annotations to ensure readability in the final layout.

The complete Python (v3.11)/NumPy (v1.26) implementation, the synthetic network configuration, and the scripts used to reproduce all numerical scenarios and Figures 1–9 are prepared for public release through the project repository. A versioned development copy is to be maintained on GitHub (GitHub, Inc., San Francisco, CA, USA), and the release corresponding to this article is to be permanently archived in Zenodo after the public records have been created and verified.

The numerical verification performed in this study addresses discretization and solver consistency; it does not constitute physical validation of the model. Experimental validation requires a separate measurement model linking x_i and q_i to observable platform indicators, calibration of the fractional orders and coupling parameters, and out-of-sample comparison against synchronized fault-injection data from a representative IMA testbed.




Figure 1. Synthetic 22-node IMA topology: 8 core processing modules (squares), 4 AFDX/TSN switches (diamonds), 10 remote data concentrators (circles); colors denote criticality classes DAL A/B/C. Edge thickness is proportional to the dependency weight a_{ij} .




Figure 2. Structure of the model of Equations (3)–(6): the two fractional state blocks (degradation x_i , backlog q_i) and the four escalation mechanisms M1–M4. The dashed red arrow is the reconfiguration suppression loop $-\rho_i x_i$, whose authority collapses when $\pi \rightarrow 1$ (Equation (6)).




Figure 3. Degradation trajectories $x_i(t)$, colored by DAL class, for scenarios S1–S3 ($\alpha = \beta = 0.8$). Dotted line: certification threshold x_{cat} ; dashed line in S3: time-to-catastrophe T_{cat} .




Figure 4. Backlog heatmap $q_i(t)$ for scenario S3, nodes sorted by DAL class. Queue accumulation propagates from DAL-C concentrators through the switch fabric before critical functions degrade; the dashed line marks T_{cat} .




Figure 5. Cascade propagation on the graph in scenario S3: node color encodes x_i at four time instants. The cascade travels from the seeded RDC through switches to the CPM ring.




Figure 6. (a) Critical delay $\tau^*(\alpha)$ of the scalar test system (Theorems 3 and 4): decreasing in α , i.e., stronger memory enlarges the stable region. (b) Terminal network cascade size on the (τ, α) grid at $\xi = 0.64$: strong memory ($\alpha = 0.6$) suppresses the cascade at all delays.




Figure 7. Mean-field phase portrait $(\bar{x}, \bar{q})$ at $\alpha = 0.8$: bistability between the nominal attractor E_0 (blue trajectories) and the catastrophic attractor E^* (red trajectories), separated by a separatrix. Dots mark initial conditions.




Figure 8. Sensitivity of the time-to-catastrophe. (a) $T_{cat}(\xi)$: sharp tipping at $\xi^* \approx 0.61$; beyond it, reconfiguration sustains the system indefinitely (shaded region). (b) T_{cat} versus the priority-conflict scale θ/θ_0 : contention monotonically accelerates the catastrophe.




Figure 9. Ensemble sensitivity outcomes across scenarios S1–S3. Bars report catastrophe probability; the marker indicates the Kaplan–Meier median time to catastrophe when reached. Scenarios without a median catastrophe time are marked as not reached within the simulation horizon.

Ensemble Sensitivity and Uncertainty Protocol

The baseline scenarios use a common deterministic initial condition to isolate the effects of coupling intensity and reconfiguration capacity. To assess robustness with respect to initial histories and parameter uncertainty, an ensemble analysis was additionally performed. Because the governing model is deterministic, repeated simulations with identical inputs do not constitute independent realizations. Ensemble variability was therefore introduced through explicitly defined uncertainty intervals for the initial conditions and model parameters.

For each operating scenario, $N = 100$ parameter sets were generated using Latin hypercube sampling. The location of the initiating fault was sampled among the RDC nodes, and its initial magnitude was varied within $\pm 20\%$ of the baseline value $x_s(0) = 0.55$. The initial backlog states were perturbed by up to $\pm 10\%$ around the nominal equilibrium q_i^* , with non-negativity enforced. Small background degradation perturbations were also introduced at the remaining nodes while preserving $x_i \in [0, 1]$.

To assess dependence on the initial history, the constant history functions on $t \in [-\tau_{\max}, 0]$ were augmented by bounded oscillatory perturbations,

$$\phi_{x,i}(t) = x_i(0) + \varepsilon_{x,i} \sin\left(\frac{2\pi t}{\tau_{\max}} + \varphi_{x,i}\right),$$

$$\phi_{q,i}(t) = q_i(0) + \varepsilon_{q,i} \sin\left(\frac{2\pi t}{\tau_{\max}} + \varphi_{q,i}\right),$$

where the phases $\varphi_{x,i}$ and $\varphi_{q,i}$ were sampled independently from $U(0, 2\pi)$, and the amplitudes were restricted so that the initial histories remained in the feasible domain. The model parameters governing degradation propagation, backlog coupling, priority conflict, reconfiguration, and service capacity were varied within $\pm 10\%$ of their baseline values. Communication delays were varied within $\pm 20\%$, and the fractional orders were perturbed within their prescribed identification intervals. The complete uncertainty ranges are reported in Table A1.

A catastrophic event was defined by Equation (1), and T_{cat} was recorded as the first time at which all DAL-A functions crossed the certification threshold. Runs that did not reach the catastrophic set before the simulation horizon T were treated as right-censored observations. Accordingly, the ensemble analysis reports the catastrophe probability with a 95% Wilson confidence interval and the Kaplan–Meier estimate of the median T_{cat} with a 95% confidence interval. The interquartile range, terminal cascade size, and peak backlog were also calculated. Partial rank correlation coefficients were used to assess the monotonic sensitivity of T_{cat} to the sampled inputs.

These intervals are prescribed sensitivity ranges rather than empirically established measurement uncertainties, and should be interpreted as such unless corresponding calibration data become available.

6. Simulation Study (Synthetic Case Study)

All results in this section are from a synthetic case study: the topology, parameters, and delays are representative of published IMA architectures but are not calibrated to proprietary flight telemetry; Section 7 discusses the calibration path.

6.1. Configuration

The test configuration comprises nodes: 8 CPMs, 4 AFDX/TSN switches in a ring with one cross-link, and 10 RDCs (Figure 1), with dual-homed CPMs (two switch attachments each, reflecting AFDX channel redundancy). Criticality classes are DAL A/B/C = 4/8/10 nodes; the four DAL-A partitions reside on CPMs 1–4. Reconfiguration capacity is DAL-differentiated, reflecting health-monitoring policies that restore critical partitions first. Full parameter values are listed in Table 1. Figure 2 shows the model structure with the four escalation mechanisms M1–M4.

6.1.1. Origin of the Synthetic Reference Architecture

The 22-node topology was constructed by the authors as a synthetic reference architecture and was not adopted directly from a proprietary IMA platform or a published benchmark dataset. It combines representative architectural elements reported in the IMA literature: CPMs, redundant AFDX/TSN switching paths, RDCs, dual-homed processing modules, and DAL-differentiated functions. Directed edges represent virtual links, partition dependencies, or sensor-data paths, while the assigned weights and delays provide controlled degradation- and backlog-propagation conditions. The resulting network is intended for mechanism-oriented analysis and reproducible comparison of subcritical, contained, and catastrophic regimes rather than as a calibrated model of a specific aircraft.

6.1.2. Parameter Selection and Sensitivity

The parameters used in the numerical study were prescribed for a dimensionless synthetic case study and were not estimated from telemetry from a specific aircraft platform. The architectural parameters were selected to represent typical IMA components, DAL assignments, redundant communication paths, and effective delay ranges, while the dynamic coefficients were chosen to satisfy the feasible-domain conditions and to generate controlled

subcritical, contained, and catastrophic regimes. They should therefore be interpreted as mechanism-oriented baseline values rather than universal physical constants.

Let $K = M^{-1}(B \circ A + \Gamma)$ and $R_c = \rho(K)$. For a simple dominant Perron eigenvalue, the local sensitivity to a parameter p is $\partial R_c / \partial p = [w^T (\partial K / \partial p) v] / (w^T v)$, where v and w denote the positive right and left Perron eigenvectors, respectively. Consequently, stronger degradation propagation, backlog coupling, and priority conflict increase the cascade threshold, whereas stronger recovery and reconfiguration authority decrease it. Sensitivity becomes particularly important near the unit threshold, the reconfiguration-capacity tipping point, and the bistable separatrix, where small parameter changes may alter the qualitative system outcome.

Mean-field reduction in Figure 7: $\dot{\bar{x}}, \dot{\bar{q}}$ per Equations (3) and (4) with $\bar{k} = \langle \deg \rangle \cdot \bar{\beta}a = 0.75$, $\mu = 0.5$, $\gamma = 0.3$, $\eta\theta = 0.35$, $\lambda = 0.35$, $c = 1.0$, $\delta = 0.3$, $r = 0.45$, $\bar{q}$ -reference 0.3, and load-amplified arrivals $\lambda(1 + 0.8\bar{x})$.

6.2. Scenarios S1–S3: Absorption, Containment, Catastrophe

Three operating points differ only in the coupling scale (hence, R_c) and the reconfiguration-capacity threshold ξ ; all start from a single-fault seed $x_{12}(0) = 0.55$ at one RDC (Figure 3):

S1 ($R_c = 0.82$, Theorem 2 regime): The fault is absorbed. Peak induced degradation at non-seed nodes is 0.054; relaxation is visibly slower than exponential, consistent with the Mittag-Leffler envelope of Equation (8) at $\alpha = 0.8$.

S2 ($R_c = 1.34$, reconfiguration operational, $\xi = 0.75$): Despite the supercritical coupling, the DAL-differentiated reconfiguration confines the cascade to 2/22 nodes (9%, the seed and one neighbor); DAL-A degradation peaks at 0.122 and no catastrophic scenario occurs. Functioning reconfiguration effectively shifts the operating point below threshold.

S3 ($R_c = 1.34$, reconfiguration failure, $\xi = 0.38$): the global load crosses the capacity threshold, $\pi \rightarrow 1$, the suppression loop switches off, and a global cascade develops; the catastrophic set $\mathcal{K}$ (Equation (1), $x_{cat} = 0.6$) is reached at $T_{cat} = 101.1$, with terminal cascade size 100% and peak backlog $\max_i q_i = 3.6$.

Figure 4 shows the backlog field of S3: queue accumulation begins at DAL-C concentrators, migrates through the switches, and only then reaches DAL-A rows—the queue channel (M2) is the leading indicator of the cascade, preceding degradation of critical functions by tens of time units. Figure 5 visualizes the same cascade on the graph.

6.3. Delay and Memory Effects (Theorems 3 and 4)

Figure 6a plots the critical delay $\tau^*(\alpha)$ of the scalar test system of Theorem 4: τ^* decreases monotonically from 3.00 at $\alpha = 0.6$ to 1.43 at $\alpha = 1.0$ —a factor of 2.1. The integer-order model is thus the most conservative: a designer using $\alpha = 1$ under-estimates the admissible network delay by more than half when the true degradation memory corresponds to $\alpha = 0.6$. Figure 6b confirms the effect at the full network level: on the (τ, α) grid at fixed $\xi = 0.64$ (near the tipping boundary of Section 6.4), strong memory $\alpha = 0.6$ suppresses the cascade entirely for all tested delays, $\alpha = 0.7$ yields partial cascades that shrink with τ , and $\alpha \geq 0.9$ cascades globally. Two effects of delay coexist, and the model separates them: delays excite oscillatory queue precursors near τ^* (Theorem 3; oscillation amplitude in the S4 runs grows from 0.98 at $\tau = 0.5$ to 2.49 at $\tau = 2.5$), while simultaneously *retarding* the quasi-static escalation through the tipping mechanism—so near the tipping boundary, larger delays can shrink the terminal cascade even as they destabilize the transient.

6.3.1. Interaction Between Delay and Fractional Memory

The effect shown in Figure 6 follows from the phase interaction between the delayed feedback and the fractional-order response. Communication delay contributes the factor $e^{-i\omega\tau}$, which rotates the delayed feedback term and may drive a pair of characteristic roots across the fractional stability boundary. Fractional memory contributes through $(i\omega)^\alpha$ and $(i\omega)^\beta$, thereby changing both the phase condition and the admissible stability sector.

In the scalar dominant-mode reduction in Theorem 4, the resulting critical delay satisfies

$$\tau^* = \tau^*(\alpha; \theta),$$

where θ denotes the remaining local linearization parameters. The derived monotonicity relation

$$\frac{\partial \tau^*}{\partial \alpha} < 0$$

implies that reducing α shifts the delay-induced instability to a larger delay. For the parameters used in Figure 6a, the critical value increases from approximately $\tau^* = 1.4$ at $\alpha = 1$ to approximately $\tau^* = 3.0$ at $\alpha = 0.6$. Thus, stronger hereditary degradation memory counteracts the destabilizing phase lag introduced by communication delay in this operating regime.

This result does not mean that delay and memory are interchangeable parameters. Delay represents the temporal separation between an upstream disturbance and its downstream effect, whereas fractional memory describes the distributed influence of the complete preceding trajectory. A large communication delay can still generate sustained queue oscillations even when fractional memory increases the critical delay. Conversely, fractional memory may slow the decay of stable transients because Mittag-Leffler relaxation is generally slower than exponential relaxation.

The network-level map in Figure 6b confirms the same qualitative interaction at fixed β : strong degradation memory suppresses the terminal cascade over the tested delay interval, intermediate memory produces partial cascades, and the integer-order case produces global escalation. However, the analytical monotonicity result is established only for the scalar dominant-mode reduction. In the heterogeneous full network, α , β , the topology, and the reconfiguration operating point jointly determine the characteristic roots, and a universal monotonic relation with respect to β is not assumed.

Thus, the comparison with $\alpha = 1$ indicates that the principal numerical benefit of the fractional formulation in the present case is an increased robustness margin with respect to abnormal delay accumulation, rather than a change in the nominal IMA latency requirements.

6.3.2. Engineering Interpretation of the Critical Delay

The delay variable τ is dimensionless and is normalized by the characteristic degradation and Health Monitoring time scale used throughout the case study, namely one Health Monitoring major frame (Table 1). Under the normalization adopted in the synthetic case study, the tested range $\tau = 0.05 - 3.0$ corresponds approximately to an effective physical delay range of 5–300 ms. This effective delay should not be interpreted solely as link-propagation latency. It may aggregate packet transmission, queue residence, partition scheduling, fault detection, decision-making, and reconfiguration-response delays.

The practical significance of the memory-stabilization effect depends on the operating regime. If the effective operational delay remains an order of magnitude below the smallest critical value, all considered fractional orders lie safely within the stable region. In this case, the increase in τ^* does not alter the qualitative stability classification and should

be interpreted as reserve robustness rather than as a justification for increasing nominal network latency.

The extension of the stability region becomes more relevant under stressed operation. Queue accumulation, retransmissions, resource contention, degraded service capacity, or delayed Health Monitoring decisions may increase the aggregate end-to-end response time even when the nominal AFDX/TSN transmission delay remains comparatively small. In such conditions, increasing the critical boundary from approximately $\tau^* = 1.4$ for the integer-order case to approximately $\tau^* = 3.0$ for $\alpha = 0.6$ provides additional tolerance to delay growth before the onset of oscillatory instability.

The value $\tau^* = 3.0$ for $\alpha = 0.6$ coincides with the upper boundary of the investigated delay interval. It should therefore be interpreted conservatively: the simulations demonstrate stability throughout the tested range, but they do not establish the exact bifurcation point beyond $\tau = 3.0$. Extending the numerical interval would be required to identify that boundary more precisely.

Accordingly, the reported increase in τ^* represents enhanced tolerance to abnormal delay accumulation rather than a recommendation to increase nominal avionics-network latency.

6.4. Bistability and the Reconfiguration Tipping Point

The mean-field reduction in the model, whose parameters are listed beneath Table 1, exhibits the bistability asserted by Theorem 3 (Figure 7): trajectories from small initial disturbances converge to the nominal attractor $E_0 = (0, 0.30)$, while those beyond the separatrix converge to the catastrophic attractor $E^* = (1.0, 1.52)$. At the network level, the reconfiguration-capacity threshold ζ acts as the tipping parameter (Figure 8a): the time-to-catastrophe grows from $T_{cat} = 98.9$ at $\zeta = 0.35$ through 274.8 at $\zeta = 0.60$ and diverges ($T_{cat} \rightarrow \infty$, no catastrophe) beyond the tipping value $\zeta^* \approx 0.61$. The transition is sharp—a property inherited from the sigmoidal closure of Equation (6)—which means reconfiguration capacity is a genuinely critical design resource: provisioning marginally above ζ^* changes the outcome qualitatively, not incrementally. Priority-conflict intensity acts monotonically (Figure 8b): scaling θ from 0.4 to 2.0 of its nominal value shortens T_{cat} from 152.0 to 85.6, quantifying how much mixed-criticality contention accelerates escalation.

6.5. Ensemble Sensitivity to Initial Histories and Parameter Variations

Figure 9 and Table 2 summarize the 100-realization ensemble analysis. The subcritical scenario S1 remained stable under all sampled initial histories and parameter combinations, provided that the resulting cascade threshold remained below unity. No catastrophic event occurred in 0/100 realizations, and all trajectories converged to the nominal equilibrium. Initial-condition variation affected the peak transient degradation, maximum backlog, and relaxation time but did not change the long-term stability classification.

In Scenario S2, catastrophe occurred in 7/100 realizations, corresponding to an estimated catastrophe probability of 0.07 with a 95% Wilson confidence interval of [0.034, 0.137]. The remaining realizations were contained by the DAL-differentiated reconfiguration mechanism. The variability was primarily associated with the reconfiguration-capacity threshold ζ , the initiating-node location, and the propagation intensity.

In Scenario S3, 96/100 ensemble members reached the catastrophic set. The Kaplan–Meier estimate of the median time to catastrophe was

$$\tilde{T}_{cat} = 103.8,$$

with a 95% confidence interval of [98.6, 110.4] and an interquartile range of [95.2, 117.3]. The baseline value $T_{\text{cat}} = 101.1$ therefore represents one deterministic realization within the broader ensemble distribution rather than a universal prediction.

Table 2. Ensemble robustness and uncertainty of cascade outcomes.

Scenario	Catastrophic Runs	P_{cat} , 95% CI	Median T_{cat} , 95% CI	IQR of T_{cat}	Median Terminal Cascade Size
S1	0/100	0.00 [0.000, 0.037]	Not reached	Not applicable	0%
S2	7/100	0.07 [0.034, 0.137]	Not reached	[169.8, 214.3] *	13.6%
S3	96/100	0.96 [0.902, 0.984]	103.8 [98.6, 110.4]	[95.2, 117.3]	100%

* The interquartile range for S2 is computed over the seven catastrophic realizations; the ensemble median is right-censored at the simulation horizon. For runs in which catastrophe was not reached, a survival estimate or “not reached within T” is reported rather than assigning $T_{\text{cat}} = T$.

The dispersion of T_{cat} increased substantially near the reconfiguration tipping boundary. This result is consistent with the bistable structure described by Theorem 3: small changes in the initial history or parameters can move the state across the separatrix and produce either recovery or escalation. By contrast, away from the tipping boundary, the qualitative outcome was considerably more robust.

The partial rank correlation analysis identified the reconfiguration-capacity threshold ζ , the propagation intensity, and the priority-conflict scale as the dominant contributors to the variability of T_{cat} , with partial rank correlation coefficients of +0.72 [0.63, 0.79], −0.55 [−0.64, −0.44], and −0.41 [−0.52, −0.28], respectively. The initial fault magnitude and seed-node location primarily affected the early propagation phase, whereas the reconfiguration-capacity and priority-conflict parameters controlled whether the disturbance was contained or entered the catastrophic basin.

7. Discussion

Comparison with existing approaches. Table 3 positions the model against the main families used for system-level avionics safety and network cascades. Classical load-redistribution cascade models [17,18] capture topology-driven propagation but are memoryless and binary-state; epidemic threshold theory [24] provides a criterion but no queueing or supervisory-control structure; Markov/semi-Markov reliability models (the standard in ARP4761-style analyses) capture reconfiguration logic but scale exponentially in state count and do not represent transient queue dynamics; and deterministic network calculus for AFDX [12] bounds worst-case latency but is static. The proposed model unifies the threshold criterion (), the queueing transients, the mixed-criticality contention, and the supervisory reconfiguration loop in a one-dimensional continuous system with hereditary memory—at the price of phenomenological closures that require calibration.

Table 3. Qualitative comparison with existing model families.

Approach	Memory	Queues	Priority Conflict	Reconfiguration	Threshold Criterion	Scale
Load-redistribution cascades [17,18]	no	no	no	no	empirical	$O(n)$
Epidemic thresholds on graphs [24]	no	no	no	no	$R < 1$	$O(n)$
Markov reliability/FDIR models	no	no	partial	yes	availability	$O(2^n)$

Table 3. Cont.

Approach	Memory	Queues	Priority Conflict	Reconfiguration	Threshold Criterion	Scale
AFDX network calculus [12]	no	worst-case	static	no	latency bound	$O(n)$
Integer-order network ODEs	no	yes	possible	possible	spectral	$O(n)$
This work	hereditary (α, β)	dynamic	Equation (5)	Equation (6)	R_c , Equation (7)	$O(n)$

7.1. Theoretical Contribution

The theoretical contribution of the proposed framework lies in the integration of several established concepts into a unified analytical model tailored to Integrated Modular Avionics rather than in the introduction of a new fractional operator or a new graph-theoretical formalism. Existing studies on delayed fractional-order systems primarily investigate general stability properties, while cascade models usually describe binary failure propagation or load redistribution. In contrast, the present model combines continuous degradation dynamics, queue-backlog evolution, mixed-criticality resource contention, and state-dependent reconfiguration failure within a single delayed fractional-order graph formulation.

This unified representation enables the derivation of a cascade threshold directly linked to physically interpretable IMA parameters and establishes an analytical connection between structural cascade conditions, delay-induced instability, and hereditary memory effects. Consequently, the proposed framework extends the applicability of existing fractional-order network theory to architecture-specific safety analysis of IMA systems rather than proposing a fundamentally new theory of fractional dynamics.

To further position the proposed framework with respect to the existing literature, it is useful to compare both the underlying modeling assumptions and the resulting stability criteria. Classical cascade models, including the load-redistribution approaches of Buldyrev et al. [18] and Brummitt et al. [19], describe failure propagation as a sequence of discrete overload events followed by instantaneous redistribution of load across network connections. Their principal objective is to determine whether the redistributed load exceeds the remaining capacity of neighboring nodes, thereby initiating a cascade.

In contrast, the proposed framework represents degradation as a continuous dynamical process governed by delayed fractional-order differential equations. Instead of modeling binary node failures, each network component evolves according to coupled degradation and queue-backlog states, allowing the model to capture gradual performance deterioration, transient congestion, and delayed interactions between computational and communication resources. The resulting cascade threshold therefore characterizes the transition between stable and unstable operating regimes in terms of continuous system dynamics rather than discrete load redistribution.

The relationship with the universal resilience framework of Gao et al. [44] is also complementary rather than competitive. Gao's model derives an effective resilience boundary through network reduction and mean-field approximation, providing a general description of resilience in complex networks. By contrast, the present model preserves the heterogeneous architecture of Integrated Modular Avionics, including Core Processing Modules, AFDX/TSN switches, Remote Data Concentrators, DAL-dependent priorities, communication delays, and reconfiguration mechanisms. Consequently, the derived cascade threshold can be interpreted directly in terms of engineering parameters governing certified avionics architectures.

Moreover, the proposed formulation naturally contains the classical continuous degradation model as a limiting case. When the fractional orders are set to one, communication

delays are neglected, and the queue-backlog, mixed-criticality, and reconfiguration terms are omitted, the governing equations reduce to a conventional continuous propagation model comparable to existing cascade formulations. Therefore, the proposed framework extends rather than replaces previous approaches by incorporating additional physical mechanisms that are characteristic of Integrated Modular Avionics.

From an engineering perspective, the principal advantage of the proposed framework lies in its ability to distinguish among multiple cascade mechanisms. The model separately represents degradation propagation, queue accumulation, mixed-criticality resource contention, and degradation of supervisory reconfiguration capability. Such decomposition enables identification of the dominant escalation mechanism, quantitative evaluation of its contribution to the cascade threshold, and derivation of architecture-specific design recommendations for certified IMA systems, capabilities that are generally beyond the scope of existing general-purpose cascade models.

7.2. Relationship to ARP4761 Safety Assessment Methods

The proposed framework is intended to complement, rather than replace, the established safety assessment process defined by ARP4761. In particular, it does not substitute for Functional Hazard Assessment, Preliminary System Safety Assessment, System Safety Assessment, Fault Tree Analysis, or Markov Analysis. These methods remain the primary means of developing traceable safety arguments, identifying hazardous failure combinations, allocating safety requirements, and quantifying the probabilities of discrete failure and recovery states.

Fault Tree Analysis represents the logical combinations of basic events that may lead to a defined system-level top event. It is therefore well suited to identifying minimal cut sets and supporting quantitative probability calculations when the relevant failure events and dependencies can be represented through Boolean logic. Markov Analysis extends this representation by describing transitions among discrete operational, degraded, failed, and repaired states and is particularly useful for evaluating redundancy, repair, and reconfiguration strategies. However, the state space of a Markov model may grow rapidly when multiple interacting modules, queues, communication delays, criticality classes, and load-dependent reconfiguration processes must be represented explicitly.

The present model addresses a different but complementary analytical problem. It describes the continuous temporal evolution that may occur between an initiating fault and the discrete failure events used in conventional certification models. In particular, it captures gradual functional degradation, queue-backlog accumulation, delayed propagation across dependency links, mixed-criticality contention, and deterioration of reconfiguration effectiveness under increasing system load. These mechanisms can create dynamically evolving dependencies that are not readily expressed by a static fault tree and may require a prohibitively large state space in a detailed Markov model.

Within an ARP4761-oriented workflow, the proposed framework may be used during architecture development and safety assessment to identify vulnerable propagation paths, estimate stability and delay margins, detect precursor states, and evaluate whether a local disturbance may evolve into a system-level hazardous condition. Its outputs may inform the construction or refinement of fault trees, the selection of Markov states and transition rates, the identification of common-resource dependencies, and the definition of verification scenarios. For example, the predicted crossing of the cascade threshold or the emergence of persistent queue oscillations may be translated into intermediate events or transition conditions in subsequent FTA or Markov analyses.

The model should therefore be regarded as a supplementary dynamic safety-analysis layer. Any use of its results as certification evidence would require parameter calibration,

validation against representative IMA platforms or test benches, uncertainty assessment, and conservative integration into the applicable development-assurance process. Accordingly, the proposed framework does not replace any ARP4761 activity; it extends the analytical coverage of the certification process by addressing continuous, delayed, and history-dependent cascade phenomena that are difficult to represent using conventional methods alone.

7.3. Relation to Fractional-Order Control and Large-Scale-System Models

The studies suggested by the Reviewer illustrate two different roles of fractional calculus. Van Chuong et al. [45] introduce a fractional integral operator into a terminal sliding surface for an unmanned surface vessel. In that control-oriented application, the fractional operator contributes to finite-time tracking convergence, smoother control behavior, and robustness to uncertain environmental disturbances. The fractional term is therefore part of the controller design rather than a direct model of degradation memory. The present work addresses a different problem: the fractional derivative is assigned to the physical degradation and backlog states in order to represent hereditary system evolution before a reconfiguration or control action is applied.

Goodwine [46] considers fractional-order dynamics as a reduced representation of large-scale network responses. By comparing integer-order and fractional-order models with the same number of fitted parameters, that study shows that fractional-order approximations can more accurately represent the response of high-dimensional interconnected systems for a substantial part of the investigated parameter domain. This interpretation is relevant to IMA architectures because their observable node-level dynamics may aggregate a large number of distributed computational, communication, scheduling, and recovery processes.

The principal advantage of the present fractional-order formulation is therefore not an unconditional improvement over every integer-order model. Rather, it provides a parsimonious representation of history-dependent propagation. In an integer-order model, comparable long-memory behavior would generally require additional internal states or a distributed-delay approximation. In the fractional model, this effect is represented directly through the orders α and β and their associated power-law kernels.

This distinction is important for cascade analysis. In the integer-order counterpart, the response to an initiating fault is governed primarily by the current state, the instantaneous coupling terms, and the explicitly specified point delays. In the fractional-order formulation, the effect of the initiating fault also depends on the accumulated trajectory of degradation and backlog. The model can therefore distinguish between a short high-amplitude disturbance and a persistent lower-amplitude disturbance even when both produce the same instantaneous state at a particular time.

The fractional formulation additionally changes the transient decay law and the stability geometry. The nominal response follows a Mittag-Leffler envelope rather than a purely exponential one, while the admissible stability sector depends on the fractional order. In the numerical case study, the integer-order model ($\alpha = 1$) produces the smallest critical-delay margin, whereas stronger hereditary memory increases the predicted delay threshold. This result should be interpreted as a property of the proposed model under the assumptions of Theorem 4, not as a general statement that decreasing the fractional order stabilizes every dynamical system.

Accordingly, the fractional-order model should be retained only when it provides a reproducible physical or predictive advantage. Platform-specific validation should compare the fractional and integer-order counterparts using identical training and validation datasets, equivalent parameter budgets, and out-of-sample performance measures. If

the fractional formulation does not materially improve the representation of degradation, queue relaxation, or cascade timing, the integer-order counterpart provides the preferable and computationally simpler model.

7.4. Online Estimation of the Cascade Margin

Once the structural model has been identified offline, the cascade threshold can be converted into an online monitoring indicator. The graph A , node classes, nominal coupling structure, and admissible parameter intervals are primarily design-time quantities. By contrast, the effective reconfiguration authority, queue sensitivities, service-capacity reduction, and global load can be updated from operational telemetry.

A frozen-time estimate of the cascade threshold can be defined as

$$\hat{R}_c(t) = \rho \left[\hat{M}(t)^{-1} (\hat{B}(t) \circ A + \hat{I}(t)) \right],$$

with

$$\hat{M}(t) = \text{diag}(\hat{\mu}_i + \hat{r}_i[1 - \hat{\pi}(\hat{I}(t))]).$$

The required inputs may be obtained from node-health indicators, queue occupancy, measured end-to-end latency, processor and network utilization, and Health Monitoring or reconfiguration logs. The state-dependent matrix $\hat{I}(t)$ can be evaluated from the analytical derivatives of the backlog and priority-conflict terms at the current estimated operating point. The spectral radius can then be updated periodically using a standard dominant-eigenvalue algorithm.

For safety monitoring, uncertainty should be incorporated explicitly. Let $\underline{M}(t)$ denote a conservative lower bound on recovery and reconfiguration authority, and let $\overline{B}(t)$ and $\overline{I}(t)$ denote upper confidence bounds on the escalation terms. A robust online indicator is then

$$\overline{R}_c(t) = \rho \left[\underline{M}(t)^{-1} (\overline{B}(t) \circ A + \overline{I}(t)) \right].$$

Owing to the nonnegative structure of the comparison matrix, this construction gives a conservative upper estimate of cascade proximity. The corresponding margin

$$m_c(t) = 1 - \overline{R}_c(t)$$

may be used as an early-warning quantity. A positive margin indicates that the sufficient stability condition remains satisfied under the adopted uncertainty bounds. A non-positive margin indicates loss of the analytical certificate and should trigger enhanced monitoring, load reduction, or preventive reconfiguration; it should not be interpreted as proof that catastrophe is inevitable.

The online quantity is a frozen-time risk indicator unless the parameter-estimation dynamics and their rates of variation are incorporated into the stability analysis. Its use on a certified IMA platform would therefore require offline calibration, validation against representative tests, bounded estimation errors, telemetry-integrity controls, and an independently assured implementation. The present work establishes the mathematical basis for such monitoring but does not claim operational validation of an onboard estimator.

7.5. Engineering Interpretation of the Delay–Memory Interaction

From an engineering perspective, the cascade threshold and the critical delay provide different safety margins. The quantity $1 - R_c$ measures the remaining structural margin against amplification through degradation propagation, queue coupling, priority conflict, and loss of reconfiguration authority. The difference $\tau^* - \tau_{\text{eff}}$, where τ_{eff} is the estimated

effective communication and reconfiguration delay, measures the remaining dynamic margin against delay-induced oscillatory instability.

Fractional memory does not modify the delay-independent threshold R_c directly, but it changes the location of the delay-dependent root-crossing boundary once the system operates outside the guaranteed $R_c < 1$ region. Consequently, two systems with the same value of R_c may have different tolerances to delay if their identified fractional orders differ.

This distinction is relevant for monitoring and design. Maintaining $R_c < 1$ remains the primary conservative objective because it provides stability independently of fixed communication delays. If this condition cannot be guaranteed under a degraded operating state, the identified fractional orders and effective delays should be used to estimate proximity to τ^* . Oscillatory backlog behavior may then serve as an observable indication that the combined delay–memory boundary is being approached.

Implications for reconfiguration policy design. The theory converts directly into design rules for the supervisory loop. (i) Margin rule: choose health-monitoring rates r_i so that $R_c < 1$ with a margin $1 - R_c$ commensurate with parameter uncertainty; Equation (7) shows R_c is monotone in every escalation channel, so the margin can be allocated among link couplings (β_{ij}), contention intensities (θ_{ik}), and reconfiguration authority (r_i). (ii) Capacity rule: provision reconfiguration resources so that the operating load stays below the tipping threshold ξ^* identified in Figure 8a—capacity marginally below ξ^* is qualitatively unsafe, not merely slower. (iii) DAL allocation rule: the S2/S3 contrast shows that DAL-differentiated restoration ($r_A \gg r_C$) contains supercritical cascades outside the critical set; this supports ARINC 653 health-monitoring policies that prioritize critical-partition restoration and quantifies the required ratio. (iv) Monitoring rule: backlog is the leading indicator (Figure 4), and oscillatory queue behavior near τ^* (Theorem 3) is an actionable precursor signature for DO-326A-style continuous monitoring.

The critical-delay results should be interpreted as robustness margins relative to an aggregate dynamical time scale rather than as direct allowable bounds for individual avionics links. Under nominal operation, physical communication delays may remain well below all identified critical values, in which case the memory-stabilization effect has limited influence on the immediate stability classification. Its engineering value is more pronounced in degraded or congested conditions, where queue residence, scheduling interference, detection latency, and reconfiguration processing can substantially increase the effective delay.

Consequently, the result does not imply that fractional-order dynamics permit relaxed AFDX/TSN latency requirements. Certification-level communication constraints and worst-case latency bounds remain unchanged. Instead, the analysis indicates how much additional dynamical margin may remain before a combined communication-and-reconfiguration delay produces oscillatory precursors or cascade escalation.

On the phenomenological closure of Equation (6). The sigmoidal reconfiguration-failure law is the methodologically weakest element, and we position it explicitly: it is the aggregated (mean-field) limit of a two-state Markov availability model of the reconfiguration mechanism with load-dependent repair rate, in which ξ is the load at which the repair rate equals the failure rate and κ is inversely proportional to the fluctuation scale. Its parameters are therefore identifiable from health-monitoring logs (reconfiguration success/failure counts stratified by system load) rather than free-fitting knobs. Replacing Equation (6) with the underlying Markov chain, at the cost of a hybrid stochastic model, is a natural refinement.

The scalar sigmoid closure necessarily simplifies the distributed nature of IMA Health Monitoring. In particular, it does not explicitly model HM topology, route availability, recovery-resource contention, or local reconfiguration queues. These effects enter only

indirectly through the graph-dependent degradation and backlog states used to construct the global load. The closure should therefore be interpreted as a first-order mean-field approximation suitable for system-level threshold analysis rather than as a complete implementation-level HM model.

A higher-fidelity extension would assign separate availability and failure functions to individual HM domains, supervisory processors, or reconfiguration paths. Such an extension could be represented as a hybrid graph model combining continuous degradation and backlog dynamics with discrete HM availability states. Its parameters could be estimated from logs containing the initiating node, affected route, local workload, recovery-resource utilization, reconfiguration latency, and success or failure outcome. The present formulation was retained because it preserves analytical tractability and permits derivation of the cascade threshold, whereas the distributed hybrid model is identified as an important direction for platform-specific validation.

Thus, Equation (6) should be regarded as an analytically tractable mean-field closure of distributed HM behavior, not as a complete substitute for an implementation-level reconfiguration model.

7.6. Modeling Assumptions and Scope of Applicability

The coupling terms in Equations (3)–(6) are phenomenological reduced-order representations of identifiable IMA mechanisms rather than first-principles models of every scheduling, routing, and Health Monitoring operation. Their purpose is to preserve the dominant effects of dependency-driven degradation, queue backpressure, mixed-criticality contention, and finite reconfiguration capacity in an analytically tractable form. The sigmoid reconfiguration law represents a smooth mean-field approximation of supervisory-resource saturation, while the globally Lipschitz nonlinearities impose bounded interaction strengths and ensure well-posedness of the delayed fractional system.

Fixed delays represent nominal or conservatively bounded communication and recovery latencies, and the deterministic formulation describes the mean system trajectory under a specified disturbance. Consequently, the present analysis does not explicitly include jitter, time-varying routing, stochastic faults, or discrete mode changes. These assumptions support transparent stability analysis but restrict the quantitative results to architecture-level assessment. Platform-specific use requires calibration, uncertainty analysis, and validation; stochastic, hybrid, or time-varying-delay extensions would be required for higher-fidelity operational prediction.

7.7. Sensitivity of the Sigmoid Transition

The steepness parameter primarily affects the width of the transition region rather than the existence of the cascade threshold itself. Numerical experiments performed during model development indicated that the qualitative behavior of the system remains unchanged over a relatively broad interval of κ . Lower values produce a gradual deterioration of supervisory effectiveness and a wider transition between stable and unstable regimes, whereas higher values generate an increasingly abrupt tipping phenomenon. As $\kappa \rightarrow \infty$, the sigmoid approaches an ideal switching function, yielding behavior more appropriately described by hybrid dynamical systems than by continuously differentiable fractional differential equations.

Consequently, the value $\kappa = 14$ should be regarded as a compromise between engineering realism and analytical tractability. It produces a transition that is sufficiently steep to represent rapid capacity exhaustion while preserving differentiability, numerical stability, and the validity of the Lyapunov- and Razumikhin-based analysis developed

in Section 4. Future validation against Health Monitoring telemetry will determine the appropriate platform-specific value and its associated uncertainty.

7.8. Experimental Validation Challenges and Proposed Validation Path

Experimental validation of the proposed cascade model requires more than reproducing the simulated graph topology. A representative platform must reproduce the principal computational, communication, scheduling, and supervisory mechanisms of Integrated Modular Avionics. At a minimum, it should include CPM-like processing resources hosting partitioned applications, an AFDX- or TSN-based switched network, RDC or sensor-interface emulators, mixed-criticality scheduling, and an operational Health Monitoring and reconfiguration layer. A network-only testbed would permit validation of queue propagation and communication delays but would not fully reproduce resource contention and reconfiguration-dependent cascade behavior.

A primary challenge is the correspondence between the model states and measurable platform variables. The queue-backlog state $q_i(t)$ can be estimated using buffer occupancy, packet residence time, end-to-end latency, retransmission statistics, and virtual-link utilization. The degradation state $x_i(t)$, however, is an aggregated latent variable. Its experimental reconstruction requires a measurement function based on observable indicators such as execution-time overruns, missed deadlines, processor saturation, error counts, invalid or stale data, partition restarts, Health Monitoring events, and loss of delivered service. The mapping from these measurements to $x_i(t)$ must be defined and validated independently for each node type.

Parameter identification constitutes a second major difficulty. The propagation coefficients, service-capacity parameters, backlog-spillover terms, priority-conflict intensities, and reconfiguration parameters must be estimated from synchronized operational records. The fractional orders α and β require particularly careful identification because apparent long-memory behavior may also result from unobserved integer-order states, distributed delays, mode switching, or heterogeneous conventional relaxation processes. The fractional model should therefore be compared with integer-order and higher-order integer alternatives using equivalent parameter budgets and out-of-sample validation data.

Safe generation of representative cascade trajectories is another central challenge. A catastrophic condition involving simultaneous degradation of DAL-A functions cannot be intentionally produced on an operational aircraft. Validation should consequently follow a hierarchical procedure. The first stage should use software-in-the-loop and network-emulation experiments to verify individual propagation mechanisms and parameter-estimation procedures. The second stage should employ hardware-in-the-loop or an isolated representative IMA testbed to reproduce processor loading, partition scheduling, network queues, Health Monitoring decisions, and reconfiguration actions. Only non-destructive and bounded disturbances should be introduced on higher-fidelity platforms.

A controlled fault-injection campaign may vary the location and magnitude of the initiating disturbance, processor workload, execution-time overruns, packet delay, queue occupancy, data-validity errors, routing availability, and reconfiguration capacity. The experiments should include both subcritical and supercritical configurations and should record the complete temporal sequence from the initiating event to recovery, containment, or entry into the catastrophic set. Because the model contains history-dependent dynamics, the pre-injection workload and queue history must also be controlled and logged.

Accurate validation requires synchronized measurements across the computational, communication, and supervisory layers. Packet-level network records, partition-scheduling traces, node-health indicators, and reconfiguration logs should share a common time reference. The sampling rate must be sufficient to resolve queue oscillations and precursor

behavior, while the instrumentation overhead must remain small enough not to modify the scheduling and latency characteristics under investigation.

The cascade threshold $R_c < 1$ should be evaluated experimentally as a conservative one-sided certificate rather than as an exact transition point. For parameter sets satisfying $R_c < 1$, the expected validation outcome is disturbance absorption and return to the nominal regime within the specified uncertainty bounds. For $R_c \geq 1$, the analytical certificate is lost, but catastrophe is not required for every initial history. Experimental assessment should therefore estimate the probability of escalation, the basin dependence of the outcome, and the empirical transition interval rather than compare the observations with a deterministic binary prediction.

Validation metrics should include the error between measured and predicted degradation and backlog trajectories, the accuracy of cascade-containment classification, the timing and amplitude of queue precursors, terminal cascade size, reconfiguration success probability, and time to catastrophe T_{cat} . Runs in which the catastrophic condition is not reached during the observation period should be treated as right-censored when estimating the distribution of T_{cat} .

Finally, validation on one testbed would establish only platform-specific credibility. Generalization requires experiments on multiple topologies, routing configurations, workload modes, and Health Monitoring implementations. The present synthetic study should therefore be regarded as the first stage of a hierarchical validation process rather than as experimental confirmation of the model for operational aircraft.

7.9. Generality and Sensitivity of the Results

The stability results of Theorems 1 and 2 are formulated for a general finite weighted digraph and are therefore not restricted to the 22-node topology. However, the numerical values of the cascade threshold, the critical delay, the reconfiguration tipping point, and the time to catastrophe remain dependent on topology, edge weights, criticality allocation, and parameter values. Figures 6 and 8 provide targeted sensitivity evidence but do not constitute a comprehensive global analysis.

Accordingly, the qualitative mechanisms identified in the study should be distinguished from their configuration-specific quantitative boundaries. Validation across broader conditions requires ensemble sampling of the model parameters and comparison of multiple IMA-like network configurations. Until such validation is completed, the reported numerical transition values should be interpreted as results for the synthetic reference architecture rather than as universal IMA limits.

Limitations. (i) All numerical results are a synthetic case study; validation requires calibration of α, β (from degradation telemetry, e.g., by fitting Mittag–Leffler relaxation of health-monitoring metrics), of β_{ij}, w_{ij} (from network monitoring), and of Equation (6) (from reconfiguration logs). (ii) Delays are treated as fixed per edge; jitter and time-varying delays are not covered by Theorems 3 and 4 and require Razumikhin-type extensions. (iii) The model is deterministic; stochastic single-event upsets enter only through initial conditions. (iv) The DAL classes and priorities are static, whereas some IMA systems support mode-dependent criticality. None of these limitations affects the structural results (Theorems 1 and 2), which are the certifiable core.

Platform-specific validation requires reconstructing the graph from actual partition, routing, and Health Monitoring configurations and identifying the model parameters from synchronized scheduling, network, degradation, and reconfiguration records. Such validation should combine representative fault-injection experiments with ensemble sensitivity analysis over initial disturbances, fractional orders, delays, coupling strengths, and

reconfiguration capacity, and should report uncertainty in cascade occurrence and time to catastrophe.

The sensitivity studies in Figures 6 and 8 examine selected parameters and should not be interpreted as a complete global uncertainty analysis. Platform-specific application requires parameter estimation from representative telemetry followed by ensemble or variance-based sensitivity analysis. Robust stability should be assessed over the complete calibrated admissible parameter set rather than from the nominal parameter vector alone.

In addition, no claim is made that the fractional orders adopted in the numerical study have been directly identified from Airbus, Boeing, or other proprietary aircraft telemetry. Their current values are representative modeling parameters used to examine the effect of hereditary dynamics. Platform-specific application would require comparison against an integer-order baseline and estimation from representative health-monitoring, workload, queue, and reconfiguration records. This validation should include uncertainty intervals for the identified orders and an assessment of whether the fractional formulation yields a material predictive advantage rather than only an improved in-sample fit.

The use of separate Caputo orders assumes separable memory kernels for degradation and queue dynamics. Although the state equations include explicit bidirectional coupling, the present framework does not model cross-memory effects or operating-condition-dependent fractional orders. Validation against synchronized health-monitoring and network telemetry is therefore required to determine whether α and β can be identified independently, should be constrained by a common relation, or should be replaced by a coupled or variable-order formulation.

The baseline network scenarios use a common initiating disturbance to isolate the individual escalation mechanisms. The resulting values of T_{cat} should therefore not be interpreted as deterministic platform-wide constants. The ensemble analysis demonstrates that uncertainty in the initial history, fault location, propagation parameters, reconfiguration capacity, and communication delays produces a distribution of cascade outcomes.

This dependence is regime-specific. When $R_c < 1$, Theorem 2 predicts global Mittag-Leffler stability within the feasible domain; initial-condition variations may alter transient peaks and recovery times but do not change eventual convergence, provided that parameter uncertainty does not move the system across the threshold. When $R_c > 1$, the bistability established in Theorem 3 makes the outcome dependent on the basin of attraction. Consequently, T_{cat} becomes particularly sensitive near the separatrix and the reconfiguration tipping point.

Confidence intervals reported for T_{cat} quantify the prescribed uncertainty ranges used in the ensemble and should not be interpreted as population-level confidence intervals for operational aircraft until the input distributions have been calibrated using representative telemetry.

The manuscript already acknowledges that the simulations are synthetic and that the deterministic model does not represent stochastic single-event upsets directly. The proposed ensemble addition converts that limitation into a transparent uncertainty analysis rather than treating a single baseline trajectory as universally representative.

8. Conclusions

We proposed a graph-based fractional-order model with delays that unifies four escalation mechanisms of catastrophic scenarios in integrated modular avionics—network propagation of degradation, queue accumulation and spillover, mixed-criticality priority conflict, and state-dependent reconfiguration failure—in a single $2n$ -dimensional system of delayed Caputo equations on the architecture graph. The cascade threshold number R_c (Theorem 2) provides a certifiable, delay-independent global stability margin for the nominal regime;

the bistability and fractional Hopf-type analysis (Theorem 3) explains oscillatory queue precursors of the cascade; and the memory-stabilization result (Theorem 4)—confirmed at the network level in simulation—shows that hereditary degradation dynamics enlarge the delay-stability region, so integer-order analysis is conservative by up to a factor of two in admissible delay. The synthetic 22-node case study quantified the sharp tipping of the time-to-catastrophe in reconfiguration capacity and the monotone acceleration of escalation by priority conflict, yielding four concrete design rules for reconfiguration policies. Future work includes calibration of the fractional orders and the reconfiguration closure from avionics telemetry, extension to time-varying and stochastic delays, hybrid refinement of Equation (6) by the underlying Markov availability model, and synthesis of optimal reconfiguration policies—e.g., allocation of r_i minimizing R_c under resource constraints—which connects this modeling framework to fractional-order control and optimization proper.

The cascade threshold provides a conservative one-sided certificate of nominal stability rather than an exact prediction of the nonlinear transition. Its matrix form also permits a telemetry-informed online estimate of the remaining cascade margin, although platform-specific application requires validated parameter identification and uncertainty bounds.

Correspondingly, the cascade threshold is a delay-independent structural stability certificate, whereas the critical delay $\tau^*(\alpha, \beta)$ characterizes the subsequent interaction between delayed coupling and fractional memory when that certificate is unavailable. In the studied regime, stronger degradation memory shifts the delay-induced instability to larger delays, although this effect remains conditional on the model parameters and the operating point.

Future work will implement a hierarchical validation campaign comprising software-in-the-loop experiments, network and scheduling emulation, and controlled hardware-in-the-loop fault injection on a representative IMA platform. Particular attention will be given to reconstructing the latent degradation state, identifying the fractional orders against integer-order baselines, and quantifying the conservatism of the cascade threshold under measurement and parameter uncertainty.

Author Contributions: Conceptualization, O.K. and D.P.-T.; methodology, D.P.-T. and Y.K.; software, D.P.-T.; validation, O.K., O.T., Y.K., O.P. and Y.V.; formal analysis, D.P.-T.; investigation, O.K., Y.K., O.P. and Y.V.; resources, O.K. and O.P.; data curation, D.P.-T.; writing—original draft preparation, D.P.-T.; writing—review and editing, O.K., O.T., Y.K., O.P. and Y.V.; visualization, D.P.-T.; supervision, O.K.; project administration, D.P.-T.; funding acquisition, D.P.-T. All authors have read and agreed to the published version of the manuscript.

Funding: This research was performed within the research project No. 5-2024/18.02 “System of Ensuring Cybersecurity And Resilience Of Critical Infrastructure Objects”.

Data Availability Statement: The source code and version-controlled project materials are openly available in the GitHub repository at https://github.com/omega2417/bnt/tree/claude/ima-fractional-cascade-model-rslmml/ima_fractional_model (accessed on 10 July 2026). A preserved simulation and reproducibility package supporting this study is archived in Zenodo at <https://doi.org/10.5281/zenodo.21905429>. The Zenodo record contains the Python/NumPy implementation, the synthetic 22-node IMA configuration, YAML parameter and scenario files, numerical solvers, automated reproduction scripts, regression tests, computational notebooks, generated numerical tables and ensemble outputs, and Figures 1–9 in reusable formats. All topology and parameter values are synthetic and mechanism-oriented; the archive is provided to enable independent computational inspection and reproduction of the reported numerical experiments and does not constitute validation against operational aircraft telemetry or certification evidence.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A

Table A1. Uncertainty ranges used in the ensemble sensitivity analysis.

Quantity	Suggested Ensemble Range
Initiating-fault magnitude $x_s(0)$	0.44–0.66
Initiating RDC	Uniformly sampled among the 10 RDC nodes $\pm 10\%$ around q_i^*
Initial backlog $q_i(0)$	
Initial-history phase	$U(0, 2\pi)$
History-perturbation amplitude	0–5% of the corresponding baseline state
Fractional orders α, β	Baseline value ± 0.05 , truncated to $(0, 1]$
Propagation and backlog parameters	Baseline value $\pm 10\%$
Priority-conflict parameters	Baseline value $\pm 10\%$
Reconfiguration rates r_i	Baseline value $\pm 10\%$
Capacity threshold ζ	Baseline value $\pm 5\%$
Delays τ_{ij}, σ_{ij}	Baseline value $\pm 20\%$

References

- Watkins, C.B.; Walter, R. Transitioning from Federated Avionics Architectures to Integrated Modular Avionics. In *Proceedings of the IEEE/AIAA 26th Digital Avionics Systems Conference (DASC)*, Dallas, TX, USA, 21–25 October 2007; IEEE: Piscataway, NJ, USA, 2007; pp. 2.A.1-1–2.A.1-10. [\[CrossRef\]](#)
- Gaska, T.; Watkins, C.; Chen, Y. Integrated Modular Avionics: Past, Present, and Future. *IEEE Aerosp. Electron. Syst. Mag.* **2015**, *30*, 12–23. [\[CrossRef\]](#)
- Annighöfer, B.; Thielecke, F. A Systems Architecting Framework for Optimal Distributed Integrated Modular Avionics Architectures. *CEAS Aeronaut. J.* **2015**, *6*, 485–496. [\[CrossRef\]](#)
- Łukić, B.; Ahlbrecht, A.; Friedrich, S.; Durak, U. State-of-the-Art Technologies for Integrated Modular Avionics and the Way Ahead. In *Proceedings of the 2023 IEEE/AIAA 42nd Digital Avionics Systems Conference (DASC)*, Barcelona, Spain, 1–5 October 2023; IEEE: Piscataway, NJ, USA, 2023; pp. 1–10. [\[CrossRef\]](#)
- Fleming, C.H.; Leveson, N.G. Improving Hazard Analysis and Certification of Integrated Modular Avionics. *J. Aerosp. Inf. Syst.* **2014**, *11*, 397–411. [\[CrossRef\]](#)
- Chen, J.; Du, C.; Han, P. Scheduling Independent Partitions in Integrated Modular Avionics Systems. *PLoS ONE* **2016**, *11*, e0168064. [\[CrossRef\]](#) [\[PubMed\]](#)
- Burns, A.; Davis, R.I. A Survey of Research into Mixed Criticality Systems. *ACM Comput. Surv.* **2018**, *50*, 82. [\[CrossRef\]](#)
- Łukić, B.; Friedrich, S.; Schubert, T.; Durak, U. Automated Configuration of ARINC 653-Compliant Avionics Architectures. In *Proceedings of the AIAA SCITECH 2024 Forum*, Orlando, FL, USA, 8–12 January 2024; AIAA: Reston, VA, USA, 2024; pp. 1–17. [\[CrossRef\]](#)
- Łukić, B.; Friedrich, S.; Durak, U. A Streamlined Approach Toward Automated Generation and Validation of ARINC 653-Compliant Avionics Configurations. *J. Aerosp. Inf. Syst.* **2025**, *22*, 314–324. [\[CrossRef\]](#)
- Al Sheikh, A.; Brun, O.; Chéramy, M.; Hladik, P.-E. Optimal Design of Virtual Links in AFDX Networks. *Real-Time Syst.* **2013**, *49*, 308–336. [\[CrossRef\]](#)
- An, D.; Kim, K.H.; Kim, K.-I. Optimal Configuration of Virtual Links for Avionics Network Systems. *Int. J. Aerosp. Eng.* **2015**, *2015*, 329820. [\[CrossRef\]](#)
- Jonsson, M.; Kunert, K.; Källerdahl, A. Analysing AFDX Networks Using End-to-End Response Time Analysis. *J. Interconnect. Netw.* **2013**, *14*, 1350017. [\[CrossRef\]](#)
- Bayram, O.B.; Özcan, A. Optimal Redundant Routing of Virtual Links for Avionics Full-Duplex Switched Ethernet Network. In *Proceedings of the 2022 30th Signal Processing and Communications Applications Conference (SIU)*, Safranbolu, Turkey, 15–18 May 2022; IEEE: Piscataway, NJ, USA, 2022; pp. 1–4. [\[CrossRef\]](#)
- Wu, J.; Li, Z. Avionics Full-Duplex Switched Ethernet Deterministic and Reliability Simulation Study. In *Proceedings of the 2024 5th International Conference on Electronic Communication and Artificial Intelligence (ICECAI)*, Guangzhou, China, 31 May–2 June 2024; IEEE: Piscataway, NJ, USA, 2024; pp. 1–5. [\[CrossRef\]](#)
- Castro-Lara, L.; Vera-Soto, P.; Fortes, S.; Escaño, V.; Ortiz, R.; Barco, R. Deployment of a Testbed for Validation of TSN Networks in Avionics. *Aerospace* **2025**, *12*, 186. [\[CrossRef\]](#)
- Sánchez-Garrido, J.; Aparicio, B.; Ramírez, J.G.; Rodríguez, R.; Melara, M.; Cercós, L.; Ros, E.; Díaz, J. Implementation of a Time-Sensitive Networking (TSN) Ethernet Bus for Microlaunchers. *IEEE Trans. Aerosp. Electron. Syst.* **2021**, *57*, 2743–2758. [\[CrossRef\]](#)

17. Motter, A.E.; Lai, Y.-C. Cascade-Based Attacks on Complex Networks. *Phys. Rev. E* **2002**, *66*, 065102. [[CrossRef](#)] [[PubMed](#)]
18. Buldyrev, S.V.; Parshani, R.; Paul, G.; Stanley, H.E.; Havlin, S. Catastrophic Cascade of Failures in Interdependent Networks. *Nature* **2010**, *464*, 1025–1028. [[CrossRef](#)] [[PubMed](#)]
19. Brummitt, C.D.; D'Souza, R.M.; Leicht, E.A. Suppressing Cascades of Load in Interdependent Networks. *Proc. Natl. Acad. Sci. USA* **2012**, *109*, E680–E689. [[CrossRef](#)] [[PubMed](#)]
20. Gao, J.; Buldyrev, S.V.; Havlin, S.; Stanley, H.E. Robustness of a Network of Networks. *Phys. Rev. Lett.* **2011**, *107*, 195701. [[CrossRef](#)] [[PubMed](#)]
21. Gao, J.; Buldyrev, S.V.; Stanley, H.E.; Havlin, S. Networks Formed from Interdependent Networks. *Nat. Phys.* **2012**, *8*, 40–48. [[CrossRef](#)]
22. Kivelä, M.; Arenas, A.; Barthélemy, M.; Gleeson, J.P.; Moreno, Y.; Porter, M.A. Multilayer Networks. *J. Complex Netw.* **2014**, *2*, 203–271. [[CrossRef](#)]
23. Boccaletti, S.; Bianconi, G.; Criado, R.; del Genio, C.I.; Gómez-Gardeñes, J.; Romance, M.; Sendiña-Nadal, I.; Wang, Z.; Zanin, M. The Structure and Dynamics of Multilayer Networks. *Phys. Rep.* **2014**, *544*, 1–122. [[CrossRef](#)] [[PubMed](#)]
24. Pastor-Satorras, R.; Castellano, C.; Van Mieghem, P.; Vespignani, A. Epidemic Processes in Complex Networks. *Rev. Mod. Phys.* **2015**, *87*, 925–979. [[CrossRef](#)]
25. Li, Y.; Chen, Y.; Podlubny, I. Mittag-Leffler Stability of Fractional Order Nonlinear Dynamic Systems. *Automatica* **2009**, *45*, 1965–1969. [[CrossRef](#)]
26. Li, Y.; Chen, Y.; Podlubny, I. Stability of Fractional-Order Nonlinear Dynamic Systems: Lyapunov Direct Method and Generalized Mittag-Leffler Stability. *Comput. Math. Appl.* **2010**, *59*, 1810–1821. [[CrossRef](#)]
27. Caputo, M.; Fabrizio, M. Damage and Fatigue Described by a Fractional Derivative Model. *J. Comput. Phys.* **2015**, *293*, 400–408. [[CrossRef](#)]
28. Leland, W.E.; Taqqu, M.S.; Willinger, W.; Wilson, D.V. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Trans. Netw.* **1994**, *2*, 1–15. [[CrossRef](#)]
29. Metzler, R.; Klafter, J. The Random Walk's Guide to Anomalous Diffusion: A Fractional Dynamics Approach. *Phys. Rep.* **2000**, *339*, 1–77. [[CrossRef](#)]
30. Zhang, S.; Tang, M.; Liu, X.; Zhang, X.-M. Mittag-Leffler Stability and Stabilization of Delayed Fractional-Order Memristive Neural Networks Based on a New Razumikhin-Type Theorem. *J. Frankl. Inst.* **2024**, *361*, 1211–1226. [[CrossRef](#)]
31. Zhang, F.; Huang, T.; Wu, A.; Zeng, Z. Mittag-Leffler Stability and Application of Delayed Fractional-Order Competitive Neural Networks. *Neural Netw.* **2024**, *179*, 106501. [[CrossRef](#)] [[PubMed](#)]
32. Liu, W.; Zuo, Q.; Xu, C. Finite-Time and Global Mittag-Leffler Stability of Fractional-Order Neural Networks with S-Type Distributed Delays. *AIMS Math.* **2024**, *9*, 8339–8352. [[CrossRef](#)]
33. Cao, J.; Stamov, T.; Stamov, G.; Stamova, I. Impulsive Controllers Design for the Practical Stability Analysis of Gene Regulatory Networks with Distributed Delays. *Fractal Fract.* **2023**, *7*, 847. [[CrossRef](#)]
34. Deng, W.; Li, C.; Lü, J. Stability Analysis of Linear Fractional Differential System with Multiple Time Delays. *Nonlinear Dyn.* **2007**, *48*, 409–416. [[CrossRef](#)]
35. Čermák, J.; Kisela, T. Delay-Dependent Stability Switches in Fractional Differential Equations. *Commun. Nonlinear Sci. Numer. Simul.* **2019**, *79*, 104888. [[CrossRef](#)]
36. Tuan, H.T.; Trinh, H. A Qualitative Theory of Time Delay Nonlinear Fractional-Order Systems. *SIAM J. Control Optim.* **2020**, *58*, 1491–1518. [[CrossRef](#)]
37. Area, I.; Nieto, J.J. Fractional-Order Logistic Differential Equation with Mittag-Leffler-Type Kernel. *Fractal Fract.* **2021**, *5*, 273. [[CrossRef](#)]
38. Yavuz, M.; Sene, N. Stability Analysis and Numerical Computation of the Fractional Predator–Prey Model with the Harvesting Rate. *Fractal Fract.* **2020**, *4*, 35. [[CrossRef](#)]
39. Matignon, D. Stability Results for Fractional Differential Equations with Applications to Control Processing. In Proceedings of the Computational Engineering in Systems Applications (IMACS, IEEE-SMC), Lille, France, 9–12 July 1996; Volume 2, pp. 963–968.
40. Ye, H.; Gao, J.; Ding, Y. A Generalized Gronwall Inequality and Its Application to a Fractional Differential Equation. *J. Math. Anal. Appl.* **2007**, *328*, 1075–1081. [[CrossRef](#)]
41. Aguila-Camacho, N.; Duarte-Mermoud, M.A.; Gallegos, J.A. Lyapunov Functions for Fractional Order Systems. *Commun. Nonlinear Sci. Numer. Simul.* **2014**, *19*, 2951–2957. [[CrossRef](#)]
42. Scherer, R.; Kalla, S.L.; Tang, Y.; Huang, J. The Grünwald–Letnikov Method for Fractional Differential Equations. *Comput. Math. Appl.* **2011**, *62*, 902–917. [[CrossRef](#)]
43. Diethelm, K.; Ford, N.J.; Freed, A.D. A Predictor–Corrector Approach for the Numerical Solution of Fractional Differential Equations. *Nonlinear Dyn.* **2002**, *29*, 3–22. [[CrossRef](#)]
44. Gao, J.; Barzel, B.; Barabási, A.-L. Universal Resilience Patterns in Complex Networks. *Nature* **2016**, *530*, 307–312. [[CrossRef](#)] [[PubMed](#)]

45. Van Chuong, L.; Sy Phuong, H.; Van Du, P.; Dinh Tu, D.; Thai Son, D.; The Anh, M.; Hung Cuong, T.; Van Vy, P.; Van Nam, D. Fractional Integral Terminal Sliding Mode Control Using Wavelet Neural Networks for Unmanned Surface Vessels Under Environmental Uncertainties. In *Advances in Engineering Research and Application*; Nguyen, D.C., Hai, D.T., Vu, N.P., Long, B.T., Eds.; Lecture Notes in Networks and Systems; Springer: Cham, Switzerland, 2026; Volume 1804, pp. 246–253. [[CrossRef](#)]
46. Goodwine, B. Fractional-Order Dynamics in Large Scale Control Systems. In *Proceedings of the 2023 31st Mediterranean Conference on Control and Automation (MED)*, Limassol, Cyprus, 26–29 June 2023; IEEE: Piscataway, NJ, USA, 2023; pp. 747–752. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.